Flexible Work, Occupational Constraints, and the Dynamics of Female Labor Supply*

Sara Casella† Kieran Larkin‡ Megha Patnaik§

July 16, 2026

Abstract

Female labor market outcomes have improved dramatically over the past half century, yet a large child penalty persists. A leading explanation is "greedy" jobs that reward long, inflexible hours, incompatible with the unequal burden of childcare on women (Goldin, 2014). We document that post-2020, these non-linear occupations disproportionately accounted for female employment growth, while hours growth has been concentrated in teleworkable occupations. We develop a dynamic general equilibrium life-cycle model with frictional occupational choice and endogenous human capital to quantify how technological shocks such as work-from-home (WFH) reshape occupational sorting and gender gaps. We discipline the WFH shock with quasi-experimental estimates of remote work's effect on maternal employment. Short-run adjustments are muted by switching frictions, but long-run general equilibrium effects are substantial: women reallocate toward non-linear occupations, female earnings rise by 12.8 percent, and the career cost of children in female earnings falls by a fifth. The welfare gain amounts to roughly a 200 percent increase in households' initial assets, reflecting lower within-household misallocation of talent and reduced consumption and hours risk. The same occupations are also the most exposed to artificial intelligence. The two technologies interact strongly: AI alone leaves the family margins largely unchanged, but combined with flexible work the demand shift implied by measured occupation-by-gender AI exposure erodes about a third of the WFH gains in the career cost of children, even as it reinforces female sorting into non-linear work and leaves the value of flexibility itself undiminished.

Keywords: Female labor supply, human capital, work from home, artificial intelligence, time use

*We thank Nick Bloom, Steve Davis, Elena Pastorino, and participants at the Stanford Remote Work Conference, Barcelona Summer Forum, FAU/IAB Macro Labor workshop (Nurnberg), the Society for Economic Dynamics 2026 (Athens), and seminars at LUISS University.

†LUISS Guido Carli and EIEF. Contact: scasella@luiss.it.

‡Institute for International Economic Studies (IIES). Contact: kieran.larkin@iies.su.se.

§LUISS Guido Carli and CEPR. Contact: mpatnaik@luiss.it.

1 Introduction

Despite decades of progress in female educational attainment and labor force participation, a substantial gender earnings gap persists in advanced economies. This gap manifests most starkly as a large "child penalty"—a sharp decline in women's earnings following the birth of their first child (Adda, Dustmann, and Stevens, 2017; Cortés and Pan, 2023; Kleven, Landais, and Leite-Mariante, 2025).[1] The leading explanation centers on the mismatch between occupational structures and household constraints. Many high-paying occupations exhibit non-linear returns to hours worked, creating what Goldin (2014) terms "greedy jobs" that reward long, continuous, and inflexible work schedules. These positions remain incompatible with childcare responsibilities that disproportionately fall on women, forcing a trade-off between career advancement and family obligations. The magnitude of this lifetime cost of occupational flexibility has profound implications for understanding persistent gender disparities in the labor market.

This paper examines how technological changes that expand workplace flexibility can reshape these occupational constraints and alter the division of labor within households. The leading case is the adoption of remote work, which weakens the link between occupational choice and a fixed workplace presence. We begin by documenting a set of new facts. Greedy occupations—those exhibiting convex returns to hours—are disproportionately *teleworkable*. Following the pandemic, remote work rose sharply and persistently, with women particularly likely to work from home on both the extensive and intensive margins. In the same period, male and female occupational outcomes have diverged: the growth in female employment shares has been concentrated precisely in teleworkable, non-linear occupations, and the joint occupational structure of couples has begun to shift. Yet switching between occupation types over the working life remains rare, especially for women. The stability of occupational sorting for incumbent workers, combined with the sudden rise in flexible work arrangements, suggests that work-from-home (WFH) may generate much larger effects in the long run than are immediately visible. These effects operate through joint household decision-making, frictional occupational mobility, and dynamic human capital accumulation—channels that require time to manifest fully.

To investigate these dynamics, we develop a heterogeneous-agent life-cycle model that quantifies how expanded workplace flexibility affects female labor market outcomes over both short and long horizons. The model features heterogeneous occupations with distinct wage-hours schedules, human capital returns that vary with work intensity, household savings decisions, and market clearing in general equilibrium. Occupations differ both in static and dynamic returns to hours worked. Couples jointly optimize consumption, savings, and labor supply subject to occupational mobility frictions, while wages adjust

[1] For example, in Denmark the arrival of a child generates a decline in female earnings of around 20 percent, and similar magnitudes are found in other developed countries (Kleven, Landais, and Søgaard, 2019; Kleven, Landais, and Leite-Mariante, 2025).

endogenously to changes in aggregate labor supply by occupation through a task-based production technology. We calibrate the model to match life-cycle patterns of work and labor supply, occupational characteristics, and gendered caregiving responsibilities in the United States, targeting empirically realistic convex returns to hours in the non-linear "greedy-jobs" sector of 1.2–1.4 (Aaronson and French, 2004).

We then use the post-pandemic expansion of telework as a natural experiment to study a large, permanent increase in workplace flexibility. In the model, remote work reduces the effective cost to the household of supplying long hours while caring for children, which we capture as a reduction in the utility cost of childcare.[2] We discipline the size of this reduction using recent quasi-experimental evidence: Harrington and Kahn (2025) for the United States and Basso et al. (2026) for Italian administrative data both estimate that access to remote work causally raises the full-year employment rate of mothers by approximately 2 percentage points, an effect very close to the aggregate change in the participation of U.S. mothers with young children between 2019 and 2024–25. Matching this short-run response in the model requires reducing the effective utility cost of childcare by 60 percent.

Our quantitative analysis yields three principal findings that highlight the importance of long-run adjustment, intra-household substitution, and general equilibrium wage effects. First, the expansion of remote work substantially reduces the career cost of children in the long run, primarily through occupational reallocation rather than hours adjustments. Female income rises by 12.8 percent while male income falls by 6.9 percent, leaving total household income slightly higher. The share of women in non-linear occupations increases by 10.9 percentage points, while the male share declines by 8.7 percentage points as households increasingly sort into configurations where women specialize in high-return greedy occupations. Women supply 7.3 percent more hours, and the employment rate of women with young children rises by 8.5 percentage points, closing around two-thirds of the motherhood employment gap. Critically, women maintain higher labor supply during peak childcare years (ages 30–40), preventing human capital depreciation and fundamentally altering life-cycle earnings trajectories. The career cost of children—the loss in the net present value of female earnings attributable to children—falls by a fifth.

Second, these effects unfold gradually due to occupational switching frictions and cohort turnover. After five years, female income rises by only 4.3 percent—a third of the long-run effect—and the female non-linear share by only 1.6 percentage points, as most incumbent workers remain in their initial occupations. The full impact materializes only as new cohorts enter the labor market and adjust their career choices from the outset. We validate the model's short-run predictions against the data: comparing the change in female relative to male outcomes between 2018–19 and 2024–25 for workers old enough to have completed their initial occupational choice, the model's five-year response accounts for essentially all

[2] This is consistent with empirical evidence from Barrero, Bloom, and Davis (2021) and Bick, Blandin, and Mertens (2022), and with gender-specific surveys on remote work preferences (Adams-Prassl et al., 2022).

of the observed relative employment change and around 60 percent of the relative shift toward non-linear occupations. The model overshoots slightly on relative full-time hours. It also captures the observed decline in relative female wages in this period, which the model attributes to a compositional effect from the entry of marginal workers.

Third, general equilibrium price adjustments matter for both the size and the distribution of the gains. The influx of women into non-linear occupations depresses the non-linear wage and raises the linear wage, a shift in relative wages of about 1.2 percent. This dampens female occupational reallocation relative to partial equilibrium (10.9 versus 12.1 percentage points), but substantially amplifies the exit of men from non-linear occupations (-8.7 versus -6.3 percentage points): as childcare becomes less specialized, households increasingly select the spouse with the higher return to specialize in the greedy occupation, and rising linear wages pull men into flexible jobs. Because women remain over-represented in linear occupations, this force reinforces the narrowing of gender income gaps. Intra-household substitution—with women specializing in greedy jobs while men move toward more flexible arrangements—represents a novel pattern of household labor allocation enabled by technological change.

We use the model to quantify the welfare consequences of flexible work. The expected value of a newly formed household rises by 1.8 percent, equivalent to roughly a 204 percent increase in average initial assets, even after accounting for general equilibrium wage adjustments. The gains reflect more than higher female earnings. A measure of within-household misallocation of talent—instances in which the spouse with the higher initial productivity draw does not occupy the high-return occupation—falls by 58 percent. The full-time female wage penalty falls by around 45 percent. And because flexibility allows households to smooth labor supply around child-related shocks, the variance of consumption growth falls by 5 percent and the variance of hours growth by 16 percent. A decomposition of the career cost of children shows that the contribution of the extensive employment margin declines, while the remaining penalty increasingly reflects slower human capital accumulation and shorter hours conditional on employment.

Finally, we turn to the next major technological disruption: artificial intelligence. Firstly, we document that the same occupations that gain from teleworking are the most exposed to AI. Secondly, exposure overall is sharply gendered—male and female non-linear occupations are similarly exposed, but the AI exposure of female-dominated linear occupations is almost double that of male linear occupations. We feed measured task-level AI exposure by occupation and gender (Eloundou et al., 2024) into the model's task-based production technology and compute the new equilibrium, comparing the WFH and AI economy to the WFH economy to isolate the interaction between the two shocks. AI acts as a further force pushing women toward non-linear occupations, because it erodes the return to the linear jobs in which women are concentrated. The two shocks interact strongly. On its own, AI leaves the family margins essentially unchanged; it is the combination with flexible work that matters. In the flexible economy, the AI demand shift undoes roughly a third of the WFH gains in the female-earnings cost of children and reduces the availability of attractive part-time options, so that mothers of young children become less likely to work. Yet the welfare value of flexibility itself is essentially undiminished in the AI economy: AI is a level shock to labor demand, not a reversal of the flexibility channel.

This paper makes three contributions. First, and most distinctively, we move the analysis of flexible work to the level at which the relevant decisions are made: the household. Occupational choice, labor supply, and childcare are joint decisions of couples, and we provide new evidence that the post-pandemic re-sorting is visible in the joint occupational structure of couples—with growth in configurations in which the wife holds the non-linear job—alongside facts linking greedy jobs, teleworkability, and AI exposure by gender. Our model is built around this margin: because spouses can substitute for one another in both market work and childcare, a fall in the cost of flexibility changes not only individual choices but the *assignment* of career intensity within the household, generating patterns of male reallocation and within-household redistribution of the career cost of children that individual-level analyses cannot capture. Second, we discipline and validate the quantitative experiment in a way that separates short-run from long-run effects: the size of the shock is pinned down by quasi-experimental estimates of remote work's effect on maternal employment, and the model's five-year response is then validated out-of-sample against the observed 2019–2025 evolution of relative female outcomes—before the model is used for the long-run questions the data cannot yet answer. The wedge between the validated short-run response and the long-run equilibrium, driven by frictional mobility, cohort turnover, and general equilibrium wage adjustment, is itself a central result. Third, the task-based production structure provides a unified framework in which work-from-home and artificial intelligence—the two major recent technological shocks to the organization of work—can be analyzed jointly, allowing us to quantify how an AI-driven demand shift interacts with the gains from flexibility rather than treating the two in isolation.

Related literature. Our work contributes to a growing literature exploring the intersection of labor market structure, family dynamics, and technological change (Albanesi and Olivetti, 2007, 2009; Greenwood et al., 2016; Erosa et al., 2022). We build in particular on Erosa et al. (2022), who study the aggregate implications of non-linear wage schedules for hours and occupational sorting, and on the analysis of the child penalty in Kleven, Landais, and Søgaard (2019), Blau and Kahn (2017), Goldin, Kerr, and Olivetti (2024), and Braun and Figueiredo (2026). Family economics has studied changes in household composition and female labor supply (Doepke and Tertilt, 2016), female participation over the life cycle (Attanasio, Low, and Sánchez-Marcos, 2008), and joint household labor supply decisions and their response to policy changes (Blundell, Pistaferri, and Saporta-Eksten, 2018; Wu and Krueger, 2021). More recently, a small literature has begun to consider the implications of labor market changes for female labor supply (Alon et al., 2020; Bang, 2022). We focus on longer-run structural changes and the implications for human capital formation. Berresheim (2025) studies a similar setting focusing on routine versus abstract task occupations. These trends have also been the focus of a broader literature on remote work flexibility (Barrero, Bloom, and Davis, 2021; Faberman, Mueller, and Sahin, 2026; Aksoy et al., 2022), including quasi-experimental evidence on remote work and maternal employment (Harrington and Kahn, 2025; Basso et al., 2026), which we use to discipline our experiment. Our AI application connects to the task-based automation framework of Acemoglu and Restrepo (2018) and to occupation-level measures of AI exposure (Felten, Raj, and Seamans, 2021; Eloundou et al., 2024; Webb, 2020).

Further, our perspective aligns with a body of work arguing that persistent gender gaps in labor market outcomes reflect structural barriers to women's entry into certain professions and educational tracks, with sizable allocative consequences for the economy (Hsieh et al., 2019; Ashraf et al., 2023). Lowering such barriers can generate sizable gains by enabling women to enter high-return occupations. Our welfare results quantify one such channel: flexibility reduces the misallocation of talent within households.

The remainder of the paper is organized as follows. Section 2 presents the motivating empirical evidence. Section 3 develops the model and Section 4 describes its calibration. Section 5 presents the quantitative results of the work-from-home experiment, and Section 6 the welfare analysis. Section 7 studies the interaction with artificial intelligence. Section 8 concludes.

2 Motivating empirical evidence

This section documents a set of facts that motivate and discipline the model. We combine data from the Current Population Survey (CPS) on occupations, hours, and employment; the American Time Use Survey (ATUS) on work location; occupation-level teleworkability from Dingel and Neiman (2020); and occupation-level artificial intelligence exposure from Felten, Raj, and Seamans (2021) and Eloundou et al. (2024).

Throughout, we operationalize the notion of "greedy" jobs following Erosa et al. (2022): we rank four-digit occupations by the average annual hours worked by male employees in the CPS and classify occupations above the median rank as *non-linear* (greedy) and those below as *linear*. Panel (a) of Figure 1 shows that this classification captures large differences in the hours margin: average hours rise steeply with the occupation's hours rank, and non-linear occupations are characterized by a high prevalence of long workweeks. This hours-based classification lines up closely with direct evidence on convex returns: Goldin (2014) and Erosa et al. (2022) document that occupations demanding long and particular hours pay disproportionately more per hour, and Aaronson and French (2004) estimate static returns to hours of 1.2–1.4 in such occupations, which we use in the calibration.

Figure 1: Greedy occupations: hours worked and employment shares by gender

(a) Hours worked and occupational hours rank

(b) Employment shares in greedy occupations

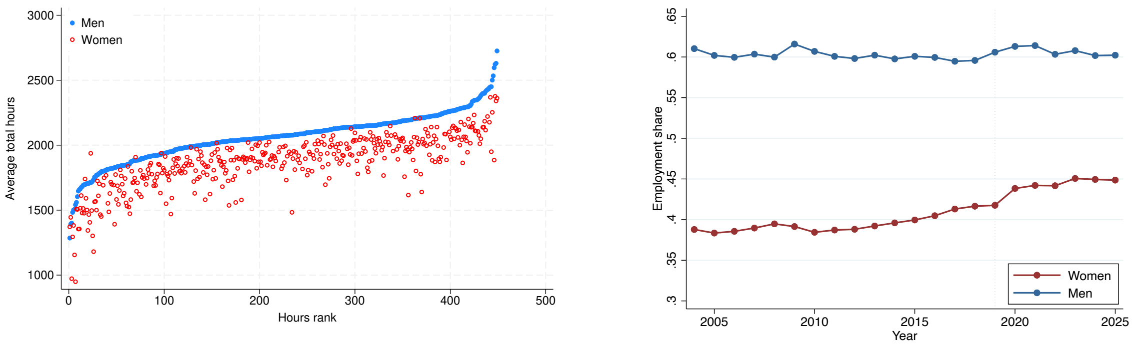

Note: Occupations (4-digit) ranked by average annual total hours for males in the Current Population Survey and classified as non-linear/greedy if they fall above the median rank, following Erosa et al. (2022). Panel (a): average annual hours by occupational hours rank. Panel (b): employment shares in greedy occupations by year and gender, CPS.

2.1 Occupational sorting by gender and the post-pandemic divergence

Women are persistently under-represented in non-linear occupations. Panel (b) of Figure 1 plots employment shares in greedy occupations by gender over time. For most of the past two decades, the gender gap in occupational sorting was large and remarkably stable. Following the pandemic, however, male and female outcomes have diverged: the share of women employed in greedy occupations has begun to rise, while the male share has not.

2.2 Greedy occupations over the life cycle and the arrival of children

The gender gap in occupational sorting is tightly linked to the family life cycle. Figure 2 plots the share of the employed working in non-linear occupations by age: the male share rises through the career, while the female share flattens and declines during the prime child-rearing ages, and Panel (b) shows the gap is concentrated among parents. Figure 3 sharpens the timing by tracing occupational status around the arrival of the eldest child in an event-study design: the arrival of a child is associated with a marked decline in mothers' attachment to non-linear occupations, with no comparable adjustment for fathers. These patterns anticipate the structure of the model, in which the career cost of children arises from the interaction of childcare with occupations that demand long hours, and is timed to the arrival of children.

2.3 The rise of remote work

Figure 4 presents trends in work location using the ATUS for full-time workers. The share of individuals working from home rose gradually over the two decades before the pandemic and then jumped discretely and persistently after 2020. Two features stand out. First, the shift is permanent: work-from-home shares remain far above their pre-pandemic trend through the most recent data. Second, women make greater use of remote work on

Figure 2: Greedy occupations over the life cycle

(a) Non-linear share by age | (b) By parental status

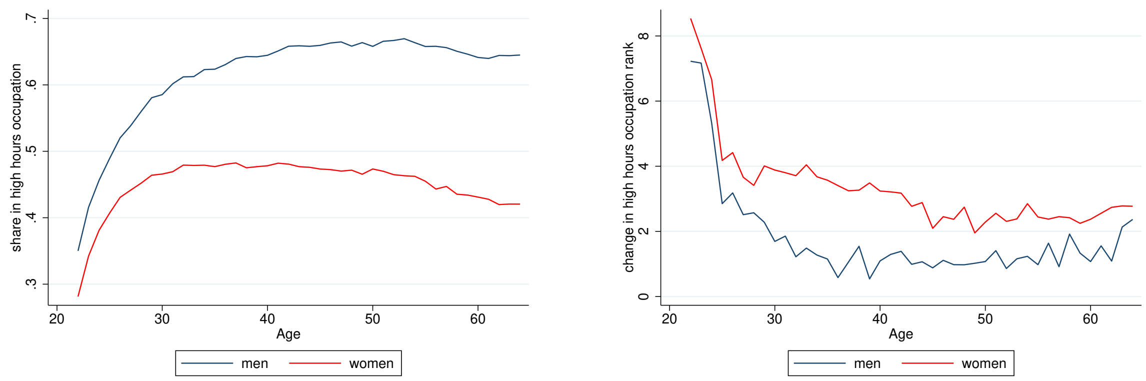

Note: CPS. Non-linear (high-hours) occupations defined as occupations for which male workers supply on average more than the median hours of work, following Erosa et al. (2022).

Figure 3: Greedy occupations and the arrival of children

(a) Non-linear share | (b) By parental status

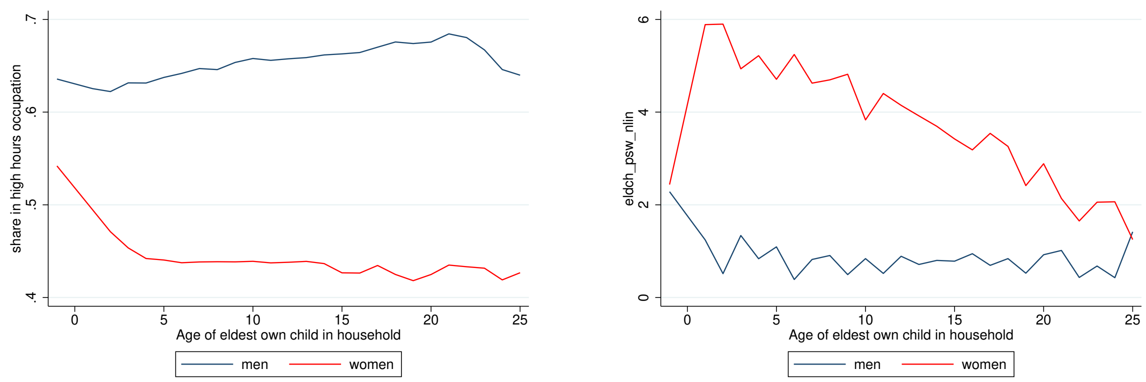

Note: CPS. Event time defined relative to the arrival of the eldest child.

both margins: they are more likely to work from home (Panel (a)), and, conditional on working from home, they supply more at-home hours (Panel (b)).

2.4 Greedy jobs are teleworkable jobs

The central empirical observation connecting these two sets of facts is that greedy occupations are disproportionately *teleworkable*. Figure 5 plots the Dingel and Neiman (2020) teleworkability measure against the occupation's hours rank: the feasibility of performing an occupation's tasks remotely rises strongly with its position in the non-linear part of the hours distribution. The correlation is robust to alternative rankings (Appendix Figure A.1). Remote work therefore relaxes the constraint that made greedy jobs incompatible with childcare exactly where returns to long hours are highest.

Is the post-pandemic reallocation of women consistent with this channel? Panel (a) of Figure 6 decomposes the post-2020 change in female employment shares into four occu-

Figure 4: Work location and hours at home, full-time employed

(a) Share working at home (%)

(b) Hours worked at home

Note: American Time Use Survey, full-time employed. Panel (a): share working "at home" on a typical day. Panel (b): hours worked at home, conditional on working from home.

Figure 5: Greedy occupations are more teleworkable

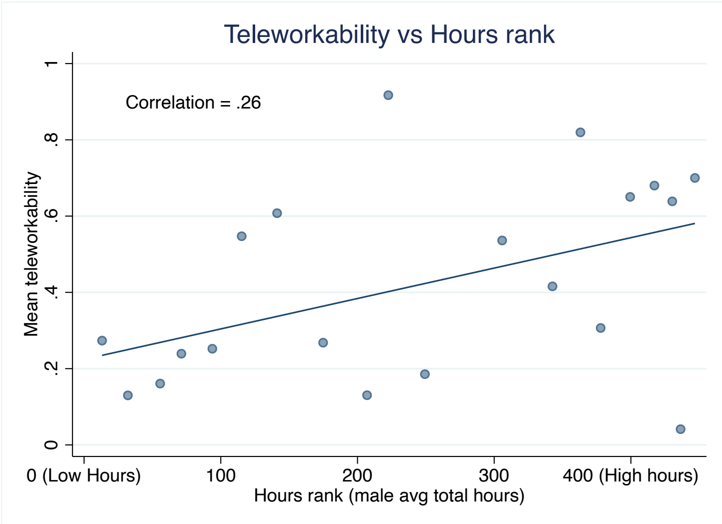

Note: Teleworkability defined following Dingel and Neiman (2020), based on the feasibility of performing tasks remotely using O*NET data. Occupations ranked by average annual hours of male workers.

pation groups defined by the interaction of the linear/non-linear and teleworkable/non-teleworkable classifications. The growth in female employment shares is concentrated almost entirely in *teleworkable non-linear* occupations; female shares in non-teleworkable linear occupations decline. Panel (b) repeats the exercise for men: no comparable pattern is present, with male employment shares changing little across the four groups. The re-sorting is thus specific to women and specific to the occupations where telework relaxed the childcare constraint on long hours.

The hours margin behaves differently, and informatively. Figure 7 plots the within-occupation growth in log hours over the same period. For women, hours growth has been strongest in the *linear* sector, in both its teleworkable and non-teleworkable segments,

Figure 6: Telework is driving the increase in female greedy-occupation employment

(a) Women (b) Men

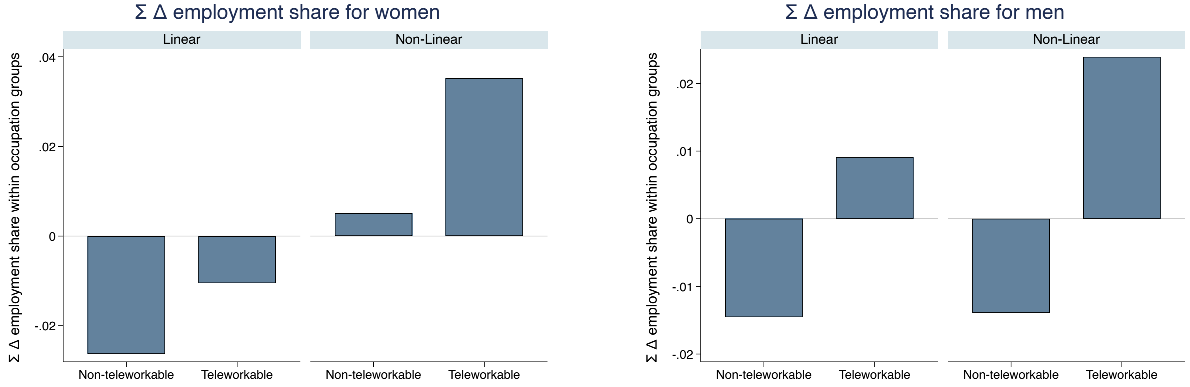

Note: Cumulative change in employment shares within occupation groups, where groups are defined by the interaction of the non-linear classification (following Erosa et al. (2022)) and the teleworkability classification (following Dingel and Neiman (2020)). CPS.

while male hours are little changed across groups. The contrast between the two margins—employment shares shifting toward teleworkable non-linear occupations while hours rise within linear occupations—is the signature of two distinct adjustments operating at different speeds: marginal workers entering employment through flexible linear jobs, and a reallocation of women toward greedy occupations that builds over time as new cohorts choose careers. The model of Section 3 reproduces both, with the first margin dominating the short run (Section 5.1) and the second the long run (Section 5.2).

Figure 7: Growth in hours within occupation groups

(a) Female hours (b) Male hours

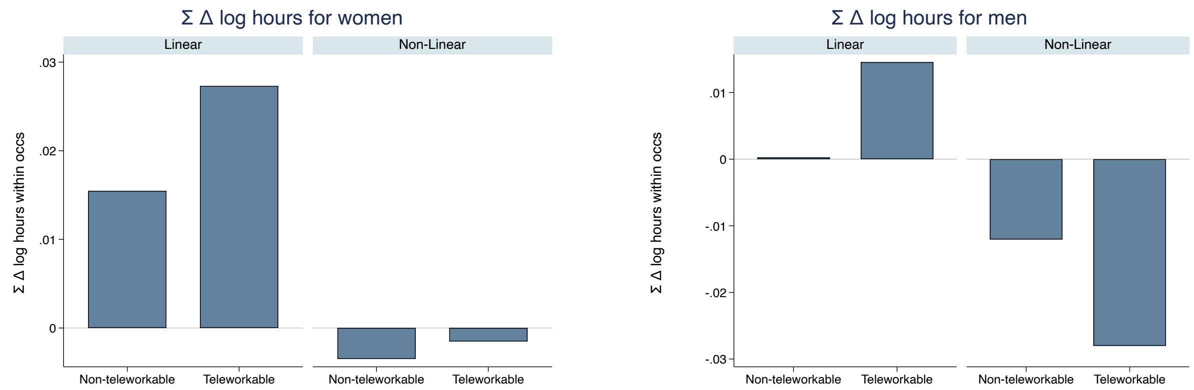

Note: Cumulative change in log hours within occupation groups defined by the interaction of the non-linear and teleworkability classifications. CPS.

2.5 Joint occupational choice at the household level

Because labor supply and childcare decisions are made jointly within households, the relevant unit for occupational sorting is the couple. Figure 8 plots the composition of couples by the joint occupational type of the two spouses over time. The traditional configuration—husband in a non-linear occupation, wife in a linear occupation—has historically been the modal arrangement. Since 2020, the joint structure has begun to shift, with growth in household types in which the wife holds the non-linear job. Appendix Figure A.3 focuses on the window around 2020. This household-level reallocation is a distinctive prediction of the mechanism we model below, in which the fall in the cost of combining long hours with childcare allows households to assign the greedy job to the spouse with the higher return rather than by gender.

Figure 8: Joint occupational choice at the household level

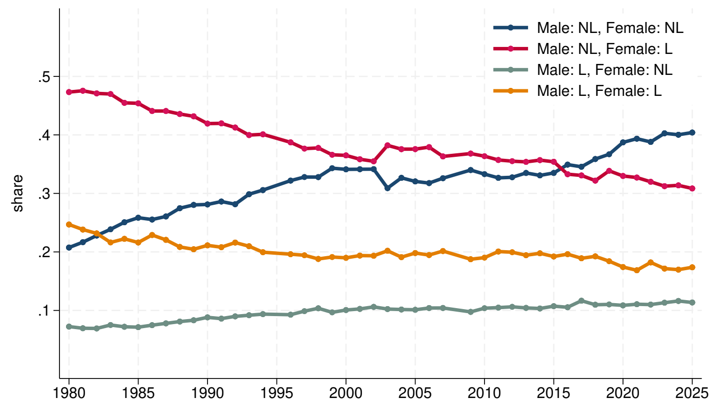

Note: Shares of married couples by the joint linear/non-linear occupational classification of the two spouses. CPS.

2.6 Occupational switching over the life cycle is rare

Finally, we document the friction at the heart of the model's dynamics. Figure 9 plots the probability of switching between linear and non-linear occupation types over the life cycle by gender. During prime working ages (25–55), occupational mobility across types is limited: the annual probability of switching is approximately 3 percent, and women exhibit a particularly low propensity to switch from linear to non-linear occupations.[3] Initial career choices therefore have persistent effects on life-cycle earnings, and aggregate adjustment to a permanent shock must operate substantially through the entry margin of new cohorts.

Taken together, the evidence points to a labor market in which (i) high-return occupations demand long hours, (ii) those same occupations have become dramatically more flexible, (iii) women—and couples—have begun to re-sort in response, but (iv) frictions imply that the re-sorting of incumbents is slow. We now build a model that embeds these ingredients and use it to quantify the short- and long-run consequences.

[3]If anything, these rates overstate true mobility, as some measured switches in the CPS reflect coding error.

Figure 9: Switches across occupation type over the life cycle by gender

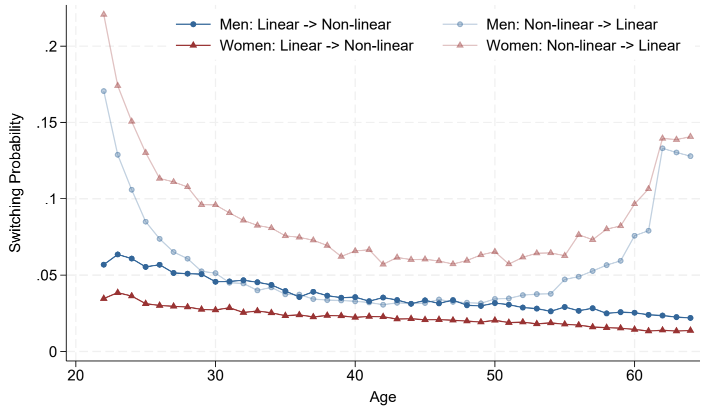

Note: Occupations (4-digit) ranked by average annual total hours for males in 2009 in the Current Population Survey and classified as linear if they fall below the median rank.

3 Life-cycle model

We develop a dynamic life-cycle model of a unitary household to analyze the interaction between occupational choice, human capital accumulation, and female labor supply. The economy is populated by households consisting of a male (m) and a female (f) who make joint decisions on consumption, savings, and labor supply. A key feature of the model is the distinction between occupation types that differ in their wage-hours schedules, capturing the "greedy" nature of certain jobs.

The life cycle is divided into three stages: an initial stage where individuals choose occupations and form couples, a working stage ($j = 1, \ldots, J^W$), and a retirement stage ($j = J^W + 1, \ldots, J^R$). Households maximize the expected discounted stream of utility, subject to budget constraints, human capital accumulation technologies, and frictional occupational mobility.

3.1 Working age

Household utility function. Households act as a unitary entity and maximize a utility function that is additively separable in consumption and the labor supply of both spouses. The per-period utility function for a household of age j is given by:

$$U(c, h_m, h_f, k, s) = k\frac{(c/k)^{1-\alpha}}{1-\alpha} - \varphi\frac{h_m^{1+\nu}}{1+\nu} - \varphi\frac{h_f^{1+\nu}}{1+\nu} - \mathcal{C}(h_m, h_f, k, s), \tag{1}$$

where c denotes total household consumption, and h_m and h_f represent hours worked by the male and female, respectively. The parameter k represents the consumption equivalence scale, which depends on the presence and age of children, and s is a childcare shock described below. The parameter α is the coefficient of relative risk aversion, ν governs the Frisch elasticity of labor supply, and φ is the disutility weight on work hours.

3.1.1 Labor supply and human capital

Individuals choose an occupation $o \in \mathcal{O} = \{\mathrm{L}, \mathrm{NL}\}$, linear or non-linear, at the start of the life cycle (age $j = 0$), and may subsequently switch subject to frictions. Occupations differ in their returns to hours worked. Labor supply is chosen from a discrete set

$$h \in \{h_{\text{nil}}, h_{\text{part}}, h_{\text{full}}, h_{\text{over}}\} = \{0,\ 0.5,\ 1,\ 1.25\},$$

corresponding to non-employment, part-time, full-time, and overtime.

Earnings function. The earnings of individual $i \in \{m, f\}$ in occupation o at age j are determined by a non-linear function of hours worked and accumulated human capital z:

$$y_o(z, h, j) = \omega_o\, \vartheta_o(h)\, \ell^o(j)\, h^{\theta_{j,o}}\, z. \tag{2}$$

Here, ω_o is the occupation-specific wage rate per efficiency unit, determined in equilibrium, and $\ell^o(j)$ captures deterministic life-cycle productivity growth specific to the occupation. The parameter $\theta_{j,o}$ governs the static return to hours: in the **linear** occupation $\theta_{j,o} = 1$ and returns are proportional to hours; in the **non-linear** ("greedy") occupation $\theta_{j,o} > 1$, so that long hours command a wage premium. In addition, part-time work in the non-linear occupation carries a proportional earnings penalty, $\vartheta_{NL}(h_{\text{part}}) < 1$, capturing the poor availability of reduced-hours arrangements in greedy jobs; $\vartheta_o(h) = 1$ otherwise. The two occupations thus differ in the entire mapping from hours to earnings, penalizing short hours and rewarding long hours in the non-linear sector.

Human capital dynamics. Productivity (human capital) z evolves stochastically and endogenously based on labor supply, reflecting a learning-by-doing mechanism. The process follows an AR(1) specification with a drift that depends on hours worked:

$$z_i' = \rho z_i + f^o(h_i) + \epsilon_i, \tag{3}$$

where ϵ_i is an idiosyncratic shock with variance σ_z^2. The function $f^o(h)$ represents the dynamic returns to work experience, which vary by occupation and are increasing in hours: non-employment generates skill depreciation, while full-time and overtime work generate positive expected wage growth, more steeply so in the non-linear occupation.[4] This mechanism implies that reducing hours not only lowers current income via $\theta_{j,o}$ and ϑ_o, but also depresses future wage growth via z'. This dual role of hours—a static premium and a dynamic return—is the key force generating a career cost of children.

[4]Numerically, the drift is implemented as hours-dependent transition probabilities across a discrete productivity grid.

Occupational switching. Occupational mobility during the working life is frictional. Opportunities to switch occupations arrive with an age-dependent probability $\lambda_o(j)$, which declines with age. Conditional on an opportunity, the household may re-optimize the occupations of both spouses, subject to a direction-specific utility cost $\delta(o, o')$ and to a proportional loss of accumulated human capital, Δ_o, consistent with the evidence on earnings losses of occupation switchers (Huckfeldt, 2022). Occupational choices are also subject to idiosyncratic taste shocks, which smooth the discrete choice.[5]

3.1.2 Family dynamics

Children. Households may enter the model with a young child (with probability 28 percent, matching the data) or have children arrive stochastically with an age-dependent probability. Once a child arrives, it transitions stochastically through developmental stages—spending three years on average in the young state and fifteen years in the older state—before leaving the household at retirement. For consumption-equivalence purposes, we track the youngest child and summarize family status by three states: no children ($k = 2$), young child ($k = 2.25$), and older child ($k = 2.5$).

Childcare costs. A central component of our analysis is the utility cost of childcare, $\mathcal{C}(\cdot)$, which depends on the hours worked by both parents, the age of the child, a childcare shock, and the flexibility of each spouse's work environment. In its general form, the cost function is

$$\mathcal{C}(h_m, h_f, k, s) = \Phi_{o_m}\, \phi_m(k)\, h_m^2 + \Phi_{o_f}\, \phi_f(k)\, h_f^2 + \max\{\Phi_{o_m}, \Phi_{o_f}\}\, \phi_c(k)\, s\, \mathbf{1}_\ell(h_m, h_f), \quad (4)$$

where $\Phi_o \leq 1$ is a shift parameter capturing the flexibility of the work environment in occupation o, and

$$\mathbf{1}_\ell(h_m, h_f) \equiv \mathbf{1}\{\, \min(h_m, h_f) \geq h_{\text{full}} \;\text{ and }\; \max(h_m, h_f) = h_{\text{over}}\}$$

indicates that both spouses work long hours: both at least full time, and at least one at overtime.

The cost is convex in each parent's hours, capturing the complementarity between childcare and time off work, and is higher when the child is young, $\phi(k{=}2.25) > \phi(k{=}2.5)$. The cost is asymmetric, weighing more heavily on mothers, $\phi_f(k) > \phi_m(k)$, which we interpret as women holding a comparative advantage in childcare production, whether for reasons of social norms or constraints. Households in which both spouses work long hours are exposed to idiosyncratic childcare shocks s, capturing the risk of arrangements breaking down when neither parent has slack time; this joint term is scaled by $\max\{\Phi_{o_m}, \Phi_{o_f}\}$, the flexibility parameter of the *less* flexible occupation, since when both spouses work long hours the arrangement is constrained by the more rigid job. In the baseline, $\Phi_o = 1$ for

[5] Analogous taste shocks apply to the discrete hours choice.

all occupations. In counterfactuals, the adoption of work-from-home technologies lowers Φ_o below one, reducing the effective cost of reconciling work and caregiving. Our main experiment imposes a uniform reduction, $\Phi_o = \Phi_{\text{WFH}}$ for all o, in which case (4) collapses to the common-scale form $\mathcal{C} = \Phi_{\text{WFH}} \left[\phi_m(k)h_m^2 + \phi_f(k)h_f^2 + \phi_c(k)\, s\, \mathbf{1}_\ell\right]$; the extension in Appendix D instead scales Φ_o with the teleworkability of each occupation, so that each spouse's childcare cost depends on the flexibility of his or her own job (see Appendix C.1).

3.1.3 Recursive formulation

In each period of the working life, the household state is $(a, z_f, z_m, o_f, o_m, k, s)$: assets, the human capital and occupation of each spouse, family status, and the childcare shock. The household chooses consumption c, savings a', and labor supply h_m, h_f. The household's problem in the working life can be written as:

$$
\begin{aligned}
V_j^{o_f,o_m}(a, z_f, z_m, k, s) = \max_{c,h_f,h_m} \Bigg\{ & k\frac{(c/k)^{1-\alpha}}{1-\alpha} - \varphi\frac{(h_m)^{1+\nu}}{1+\nu} - \varphi\frac{(h_f)^{1+\nu}}{1+\nu} - \mathcal{C}(h_f, h_m, k, s) \\
& + \beta\, E_j \Bigg[(1-\lambda_o(j))\, V_{j+1}^{o_f,o_m}(a', z_f', z_m', k', s') \\
& + \lambda_o(j) \max_{o_f', o_m'} \left\{ V_{j+1}^{o_f',o_m'}(a', z_f', z_m', k', s') - \delta(o, o') \right\} \Bigg] \Bigg\} \qquad (5)
\end{aligned}
$$

subject to:

$$
\begin{aligned}
c + a' &= aR + y^{o_f}(z_f, h_f, j) + y^{o_m}(z_m, h_m, j) && (6)\\
z_f' &= \rho z_f + f^{o_f}(h_f) + \epsilon_f && (7)\\
z_m' &= \rho z_m + f^{o_m}(h_m) + \epsilon_m && (8)\\
a' &\geq \underline{a} && (9)
\end{aligned}
$$

where the expectation is over productivity, fertility, and childcare shocks, and, upon a switch, human capital falls by the proportion Δ_o.

3.2 Initial career choice

We assume that individuals make an initial career choice at age $j = 0$, prior to forming a household and entering the labor market. This decision is based on the maximization of expected lifetime utility, taking into account the future value of marriage and labor market outcomes.

Let $o_x \in \mathcal{O}$ denote the occupation choice for gender $x \in \{m, f\}$. Individuals face occupation-specific entry costs, denoted by $\Omega(o_x)$, which capture barriers to entry such as educational requirements or training costs, and vary across occupations as well as genders. In addition to these structural costs, individuals are subject to idiosyncratic taste shocks for each occupation, distributed according to a Gumbel distribution with scale σ_o.

The value of choosing occupation o_x is determined by the expected value of the household formed in the first period of the working life ($j = 1$). Since matching in the marriage market is assortative with respect to occupation, individuals form expectations over the occupation of their future spouse. Let $\Pr(o_{x'}|o_x)$ be the probability of matching with a spouse of occupation $o_{x'}$ conditional on one's own choice o_x. This probability distribution captures the degree of assortative mating observed in the data.

Formally, the problem for an individual of gender x at age $j = 0$ is:

$$V_x = \max_{o_x \in \mathcal{O}} \left\{ \sum_{o_{x'}} \left[\hat{V}^{o_x, o_{x'}} - \Omega(o_x) \right] \Pr(o_{x'}|o_x) + \epsilon_{o_x} \right\}, \tag{10}$$

where $\hat{V}^{o_x, o_{x'}}$ represents the expected value of a household with occupational composition $(o_x, o_{x'})$ at the start of the working life:

$$\hat{V}^{o_x, o_{x'}} = \int V_1(a_0, z_{f,0}, z_{m,0}, k_0, s_0) \, dF(a_0, z_{m,0}, z_{f,0}, k_0, s_0). \tag{11}$$

The integral is taken over the initial distribution of assets, productivities, family status, and childcare shocks, and ϵ_{o_x} is the taste shock. In the baseline, we take the initial joint occupational distribution from the data and estimate the entry costs $\Omega(o_x)$ that rationalize it; the conditional matching probabilities $\Pr(o_{x'}|o_x)$ capture the observed assortative matching, while the marginal distributions $\Pr(o_x)$ are equilibrium objects that adjust in counterfactuals.

3.3 Retirement

At age $J^W + 1$, both individuals retire simultaneously and live for a maximum of J^R periods. During this phase, households face a simplified consumption-savings problem subject to mortality risk. Labor income is replaced by a pension system that depends on the permanent income accumulated during the working life, $y = \omega_{o_m} \ell^{o_m}(J^W) z_m + \omega_{o_f} \ell^{o_f}(J^W) z_f$.

We assume a joint pension system with a replacement rate ζ. Households face individual mortality risk ξ_j, which is age-dependent but not gender-specific. The household state is augmented by a demographic variable $d \in \{1, 2\}$, indicating whether the household comprises a couple or a single survivor. If one spouse dies, the household transitions to a single household ($d = 1$) and receives a survivor's pension equal to a fraction $\hat{\zeta}$ of the joint pension. If the single member dies (or both members of a couple die simultaneously), the household dissolves and leaves a bequest, from which it derives utility via a "warm-glow" motive $B(a')$.

The value function for a retired household at age $j > J^W$ is given by:

$$W_j(a,y,d) = \max_{c,a'} \left\{ \frac{(c/d)^{1-\alpha}}{1-\alpha} + \beta \Big[\xi_j^d W_{j+1}(a',y,d) + \mathbb{I}[d=2]\, 2\xi_j(1-\xi_j) W_{j+1}(a',y,d-1) + (1-\xi_j^d) B(a') \Big] \right\} \tag{12}$$

subject to the budget constraint

$$c + a' = aR + \left[\zeta \cdot \mathbb{I}[d=2] + \hat{\zeta} \cdot \mathbb{I}[d=1] \right] y \tag{13}$$

and the borrowing constraint $a' \geq \underline{a}$. The term ξ_j^d represents the probability that the household structure remains unchanged (both survive, or the single survivor survives). The second term in the expectation captures the transition from a couple to a single household due to the death of one spouse. The final term captures the utility from bequests in the event of household dissolution.

3.4 Aggregate production

Labor markets are competitive. Output is produced by combining capital and the labor of the two occupations through a continuum of tasks, following the framework of Acemoglu and Restrepo (2018). This task structure serves two purposes: absent automation it delivers a standard CES aggregate across occupations, so that relative wages respond endogenously to occupational labor supplies; and it provides a natural laboratory for the artificial intelligence experiment of Section 7, in which AI automates a subset of tasks.

Workers in the non-linear occupation supply skilled labor services S and workers in the linear occupation supply unskilled services U. Within an occupation, male and female efficiency units are perfect substitutes:

$$S_t = \sum_{g \in \{m,f\}} \phi_{NL}^g \, e_{NL}^g, \qquad U_t = \sum_{g \in \{m,f\}} \phi_L^g \, e_L^g, \tag{14}$$

where e_o^g denotes the aggregate efficiency units supplied by gender g in occupation o, integrating $y_o(z,h,j)/\omega_o$ over the distribution of households, and ϕ_o^g are group-specific productivities, equal to one in the absence of AI.

Tasks $x \in [0,1]$ are allocated across capital, unskilled labor, and skilled labor by comparative advantage, with capital performing tasks $[0, I]$, unskilled labor $(I, \tilde{x}]$, and skilled labor $(\tilde{x}, 1]$. Aggregating across tasks yields the production function

$$Y_t = A_t \Big[m_{K,t} \, (A_K K_t)^{\zeta_c} + m_{U,t} \, (A_U U_t)^{\zeta_c} + m_{S,t} \, (A_S S_t)^{\zeta_c} \Big]^{1/\zeta_c}, \tag{15}$$

where the shares reflect the task ranges: absent automation, $m_{K,t} = I_t$, $m_{U,t} = \tilde{x}_t - I_t$, and $m_{S,t} = 1 - \tilde{x}_t$, while automation reassigns the measure of automated tasks from labor to capital (Appendix B). We parameterize the comparative-advantage schedules so that the local elasticity of substitution between the two occupations equals ϵ_L, disciplined by evidence on the skill premium (Autor, Katz, and Kearney, 2008). In the work-from-home experiments no tasks are automated, thresholds are fixed, and (15) reduces to a three-factor CES: a large inflow of workers into non-linear occupations then depresses the non-linear wage per efficiency unit ω_{NL} relative to ω_L. In the AI experiment, automation reallocates tasks from labor to capital in proportion to measured exposure by occupation and gender, and the associated cost saving raises aggregate productivity (Acemoglu, 2025). The rental rate of capital is held at its baseline value, with the capital stock adjusting. Appendix B details the task aggregation, the treatment of gender, and the mapping from measured AI exposure into the model.[6]

Equilibrium. A stationary equilibrium consists of household value and policy functions, occupation-level wages per efficiency unit $\{\omega_L, \omega_{NL}\}$, initial occupational distributions, and a stationary distribution of households over states such that: households optimize; the initial occupational choice problem is consistent with the value functions and matching probabilities; wages equal the marginal products implied by (15); and the distribution reproduces itself. In counterfactuals we compute both the new stationary equilibrium and the transition path following the shock.

4 Calibration

The model is solved and calibrated to match key features of labor supply, human capital accumulation, and occupational sorting observed in U.S. microdata, including the Current Population Survey (CPS), the American Time Use Survey (ATUS), and the Panel Study of Income Dynamics (PSID). Our calibration strategy proceeds in two steps. First, we fix a set of parameters based on standard values established in the macro-labor literature or estimated directly from empirical evidence independent of our model structure. Second, the remaining structural parameters, primarily those governing preferences and technology, are estimated by simulated method of moments, targeting a set of key moments describing household labor supply and occupational sorting patterns.

4.1 Externally calibrated parameters

Table 1 summarizes the parameters set *a priori*.

Preferences and demographics: We set the time discount factor β to 0.96 and the inverse of the elasticity of intertemporal substitution, α, to 2, consistent with standard values used in heterogeneous-agent models. The curvature of the disutility of hours worked, ν, is set

[6] Results are similar under a capital-skill complementarity specification following Krusell et al. (2000).

to 0.5. The working life spans $J^W = 40$ periods (ages 25–64), and the retirement period $J^R = 25$ periods (ages 65–90). The gross interest rate R is fixed at 1.04. Mortality risk ξ_j is taken from Social Security actuarial life tables, and the joint pension replacement rate ζ is set to 0.4, with the single-survivor rate $\hat{\zeta}$ at 75% of the joint rate. Fertility probabilities are taken directly from the CPS fertility supplement to match the proportion of households with a child by age.

Human capital and production: The productivity process z is parameterized using evidence on wage dynamics. The persistence parameter ρ is set to 0.92, following Carter et al. (2025), and the innovation dispersion σ_z is set to 0.2. The occupation-specific returns to hours worked, $\theta_{j,o}$, are fixed at $\{1, 1.2\text{–}1.4\}$ to reflect the non-linear returns in the high-return occupation, based on Aaronson and French (2004). The hours-dependent drift of human capital, $P(z', h)/\Delta$, which captures expected wage growth as a function of labor supply, is set to match PSID estimates of wage losses from non-employment and wage returns to experience, separately by occupation. The proportional productivity loss upon an occupational switch is set to 15 percent following Huckfeldt (2022). Finally, the elasticity of substitution between non-linear and linear occupations, ϵ_L, is fixed at 1.6 based on the evolution of the college premium (Autor, Katz, and Kearney, 2008).

Table 1: Externally calibrated parameters

Parameter	Symbol	Value	Source
Preferences			
Discount factor	β	0.96	standard
1/EIS	α	2	standard
Hours utility curvature	ν	0.5	standard
Earnings and occupations			
Non-linear pay parameter	$\theta_{j,o}$	$\{1, 1.2\text{–}1.4\}$	Aaronson and French (2004)
Prod. persistence	ρ	0.92	Carter et al. (2025)
Prod. dispersion	σ_z	0.2	standard
Prod. drift, no work (NL)	$P(z', h = 0)/\Delta$	-0.025	PSID wage loss
Prod. drift, part time (NL)	$P(z', h = 0.5)/\Delta$	0.02	PSID wage return
Prod. drift, full time (NL)	$P(z', h = 1.0)/\Delta$	0.03	PSID wage return
Prod. drift, no work (L)	$P(z', h = 0)/\Delta$	-0.01	PSID wage loss
Prod. loss on occupation switch	Δ_o	0.15	Huckfeldt (2022)
Elasticity of NL–L substitution	ϵ_L	1.6	Autor, Katz, and Kearney (2008)
Demographics and retirement			
Working-age periods	J^W	40	ages 25–64
Retirement periods	J^R	25	ages 65–90
Fertility probability	$\rho^k(j)$	–	CPS fertility supplement
Pension replacement rate	ζ	0.4	
Survivor pension	$\hat{\zeta}$	$0.75 \times \zeta$	
Mortality risk	ξ_j	–	SSA actuarial life tables
Gross interest rate	R	1.04	standard

4.2 Internal calibration via simulated moments

The remaining structural parameters, which are central to the mechanism of occupational choice and the child penalty, are estimated by simulated method of moments. Table 2 lists the parameters, their estimated values, and the moments that most directly identify them.

The disutility of hours φ, the part-time penalty in the non-linear occupation ϑ_{NL}, and the overtime human capital drift $P(z', h = 1.25)/\Delta$ are identified by the distribution of male and female workers across part-time, full-time, and overtime status, together with non-linear wage growth. The parameters of the childcare cost function—$\phi^m(k)$, $\phi^f(k)$, and the shock loading $\phi^c(k)$, separately for young and older children—are identified by the gaps in employment rates of mothers and fathers with young and older children relative to their childless counterparts, and by the change in the share of households in which both spouses work overtime upon the arrival of a child. The estimation attributes essentially the entire childcare burden to mothers ($\phi^m \approx 0$), with the cost of a young child more than five times that of an older child. The deterministic life-cycle productivity component, $\ell^o(j) = \mu_o + \gamma_o^1 \,\text{age} + \gamma_o^2 \,\text{age}^2$, and the relative non-linear wage ω_{NL} are identified by male wage growth between ages 25 and 50 in each occupation and by the male non-linear wage premium. The arrival rate of switching opportunities $\lambda_o(j)$ and the direction-specific switching costs δ are identified by the frequency of occupational switches over the life cycle documented in Figure 9, while the entry costs $\Omega(o_x)$ are chosen to reproduce the initial occupational distribution of men and women. Finally, the dispersion of the initial occupational taste shock, σ_o, is disciplined by the quasi-experimental estimate in Traiberman (2019) that a 1 percent increase in an occupation's wage raises the probability of choosing that occupation by 4.4 percentage points; Appendix C.2 details the implementation of this target in the model.

Table 2: Internally calibrated parameters

Parameter	**Symbol**	**Value**	**Main target**
Disutility of hours	φ	4.11	PT/FT/OT shares by gender
Part-time penalty (NL)	ϑ_{NL}	0.23	Female part-time share
Prod. drift, overtime (NL)	$P(z', h = 1.25)/\Delta$	0.025	Non-linear wage growth
Childcare, young child	ϕ^f, ϕ^m, ϕ^c	2.27, 0.00, 0.38	Emp. gaps, young child
Childcare, older child	ϕ^f, ϕ^m, ϕ^c	0.41, 0.00, 0.02	Emp. gaps, older child
Relative non-linear wage	ω_{NL}	1.31	Male NL wage premium
Life-cycle profile (NL)	$\gamma_{NL}^1, \gamma_{NL}^2$	0.0125, -0.0001	Wage growth 25–50 by occ.
Switching opportunity	$\lambda_o(1)$	0.20	Life-cycle switching rates
Switching cost	$\delta_{\to L}, \delta_{\to NL}$	0.54, 0.07	Switching rates by direction
Initial occ. taste dispersion	σ_o	1.30	Traiberman (2019) elasticity
Entry costs	$\Omega(o_x)$	–	Initial occ. distribution

Note: $\lambda_o(j)$ declines with age from its entry value. The Traiberman (2019) target is that a 1 percent increase in an occupation's wage raises the probability of choosing it by 4.4 percentage points.

4.3 Targeted moments and model fit

Table 3 compares the targeted moments in the data with those generated by the calibrated model. The model performs well on the moments most closely tied to its central mechanism. It reproduces the steep non-linear and flat linear wage-growth profiles, the large negative employment gap of mothers with young children, and the near-absence of an employment response of fathers—the gender-by-family-status differentials that identify the childcare cost function and drive the career cost of children. The fit is weaker on the composition of hours conditional on working: the coarse discrete hours grid, with a single full-time and a single overtime point, shifts mass from the full-time bin toward part-time and overtime for both genders, so that individual bin shares are matched less precisely than the participation margin itself. The model also generates a smaller employment gap for mothers of older children than in the data (-0.6 against -7.6 percent), concentrating the child penalty in the young-child years. Because the older-child state is long-lasting, this margin implies that the model, if anything, understates the lifetime career cost of children, making our estimates of the gains from flexibility conservative along this dimension.

Table 3: Targeted moments: data vs. model

Calibration moment	**Data (%)**	**Model (%)**
Male		
Part-time	6.2	9.1
Full-time	49.6	29.6
Over-time	27.3	29.3
Δ Emp. rate (<45), young child	4.2	7.0
Δ Emp. rate (<45), old child	3.2	4.7
Non-linear wage growth, age 25 to 50	84.1	66.4
Linear wage growth, age 25 to 50	31.7	30.1
Female		
Part-time	14.7	22.5
Full-time	35.3	23.7
Over-time	6.9	13.4
Δ Emp. rate (<45), young child	-20.5	-15.4
Δ Emp. rate (<45), old child	-7.6	-0.6

Note: Employment-rate moments are gaps relative to comparable adults without children, for household heads younger than 45. Wage growth is cumulative between ages 25 and 50 for men.

Untargeted life-cycle profiles. To assess the model's performance beyond the targeted moments, Figure 10 compares the model-simulated life-cycle profiles of occupational sorting and earnings against the corresponding data moments from the CPS. The left panel plots the share of the employed population working in non-linear occupations by age and gender. The model successfully reproduces the large and persistent gender gap observed in the data: it captures the hump-shaped profile of male participation in non-linear sectors and, crucially, the markedly lower and flatter female profile during prime child-rearing years, consistent with the hypothesis that family constraints deter women from entering or staying in high-intensity occupations. The right panel displays life-cycle income profiles. The model generates the steep income growth trajectory for men, driven by the dynamic returns to hours worked and the static premium in non-linear jobs, and the much flatter profile for women, resulting from a combination of lower hours worked, part-time penalties, and selection into linear occupations with lower wage growth potential. Table 4 assesses the model against untargeted moments. The model closely matches non-linear occupation shares by age for men, the direction and relative magnitude of directional switching rates, and the non-linear wage premium, and it comes close on female employment at older ages. Two limitations are worth noting. The model understates male employment at ages 45–55, reflecting income effects on asset-rich households late in the working life under the coarse hours grid; since the experiments operate primarily through the family margin at younger ages, this affects levels more than the responses we study. It also understates the (likely mismeasured) gross switching rates in the CPS, and generates a positive rather than negative change in the share of households in which both spouses work overtime upon the arrival of an older child, though it captures the large negative response for young children that identifies the joint childcare shock.

Table 4: Additional untargeted moments: data vs. model

	Data	Model
Male		
Employment rate (45–55)	86.7	66.3
Non-linear share ($<$ 45)	62.5	66.5
Non-linear share (45–55)	68.6	68.1
% Switching occupation: NL $\rightarrow$ L	4.0	0.0
% Switching occupation: L $\rightarrow$ NL	4.0	1.2
Non-linear wage premium	31.7	24.1
Female		
Employment rate (45–55)	62.1	57.8
Non-linear share ($<$ 45)	42.9	33.6
Non-linear share (45–55)	41.9	35.6
% Switching occupation: NL $\rightarrow$ L	7.0	1.2
% Switching occupation: L $\rightarrow$ NL	2.5	0.6
Δ households both working OT, young child	-6.0	-10.0
Δ households both working OT, old child	-2.7	5.0

Note: Untargeted moments. Non-linear shares are conditional on employment. Switching rates are annual and, in the CPS, likely overstated by occupational coding error. The final two rows report the change in the share of households in which both spouses work overtime, upon the arrival of a child.

Figure 10: Model fit: life-cycle profiles

(a) Non-linear occupation share (b) Income by gender

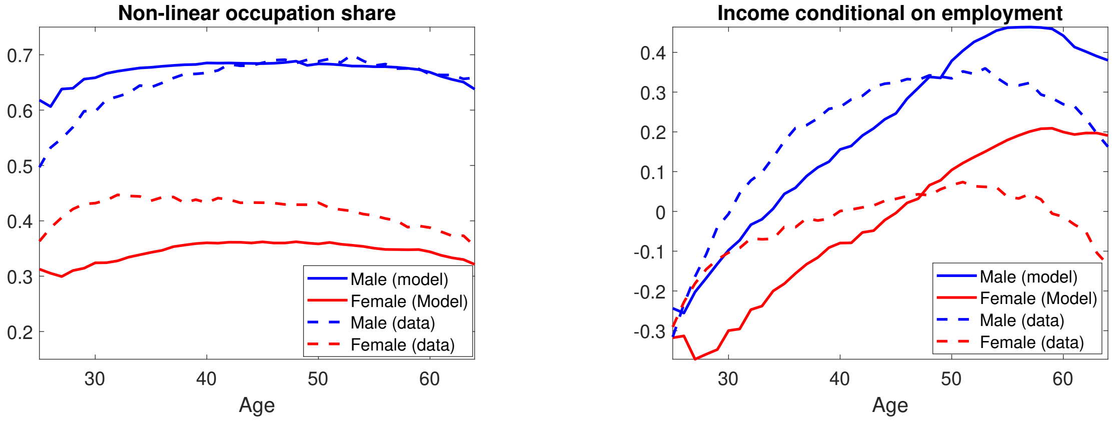

Note: Model-simulated profiles against CPS data. Left: share of the employed in non-linear occupations by age and gender. Right: mean income by age and gender.

4.4 Measuring the work-from-home shock

The effectiveness of WFH is captured by a reduction in the childcare cost scale, Φ. We discipline the size of the reduction with recent quasi-experimental evidence on the causal effect of remote work on maternal employment. Harrington and Kahn (2025) exploit variation in occupational teleworkability in the United States and estimate that the post-pandemic expansion of remote work raised mothers' full-year employment by approximately 2 percentage points. Basso et al. (2026) provide complementary evidence from Italian matched employer–employee administrative data, comparing mothers in jobs where remote work became feasible after the pandemic to mothers in comparable jobs where it did not. Their event-study estimates, reproduced in Panel (b) of Figure 11, show a persistent post-2020 increase in mothers' employment of approximately 2 percentage points in remote-feasible jobs, closing a substantial part of the motherhood employment penalty. These causal estimates are close to the raw aggregate change: in the CPS, the labor force participation of mothers with young children rose by 2.3 percentage points between 2019 and 2023–24, with no comparable change for fathers (Panel (a) of Figure 11).

In the model, we target the *change* in the employment gap, conditional on age, between women with and without a young child, in the first year after the shock, holding fixed the baseline distribution of matches, occupations, and assets. The model's baseline motherhood employment gap is -11.9 percentage points; generating a 2 percentage point short-run reduction requires lowering the childcare cost scale to $\Phi_{\text{WFH}} = 0.4$, i.e., a 60 percent reduction in the effective utility cost of childcare. The structural model then delivers the medium-run transition and the long-run general equilibrium consequences of this measured change in flexibility. Appendix C.1 details the calibration procedure and an alternative implementation of the target that mirrors the event-study designs of the quasi-experimental studies.

Discussion: flexibility and productivity. Our experiment models remote work as a pure reduction in the utility cost of combining work and childcare, holding the earnings technology fixed. A growing empirical literature suggests that remote work may also affect output per hour, mentoring, promotion, and on-the-job learning, particularly for younger workers (Emanuel, Harrington, and Pallais, 2026; Gibbs, Mengel, and Siemroth, 2023; Atkin, Schoar, and Shinde, 2023). In our framework such effects would enter the earnings function $y_o(\cdot)$ or the human capital drift $f^o(h)$, and could attenuate the human-capital-preservation channel that is central to our long-run results—or reinforce it, if flexibility primarily raises the effective labor supply of constrained workers. Two considerations limit the scope of this concern for our headline experiment. First, our calibration is disciplined by the *observed* employment response of mothers to remote-work access, which already embeds any equilibrium productivity consequences present in the quasi-experimental estimates. Second, the long-run reallocation operates through the entry margin of new cohorts choosing careers under the new technology, for whom the relevant comparison is across occupations rather than across work locations within an occupation. Nevertheless, embedding an explicit productivity margin of remote work—allowing the technology to shift both the cost of flexibility and the return to experience—is an interesting extension that we leave for future work.

Figure 11: Employment of mothers and the expansion of remote work

(a) United States, CPS

(b) Italy, administrative data

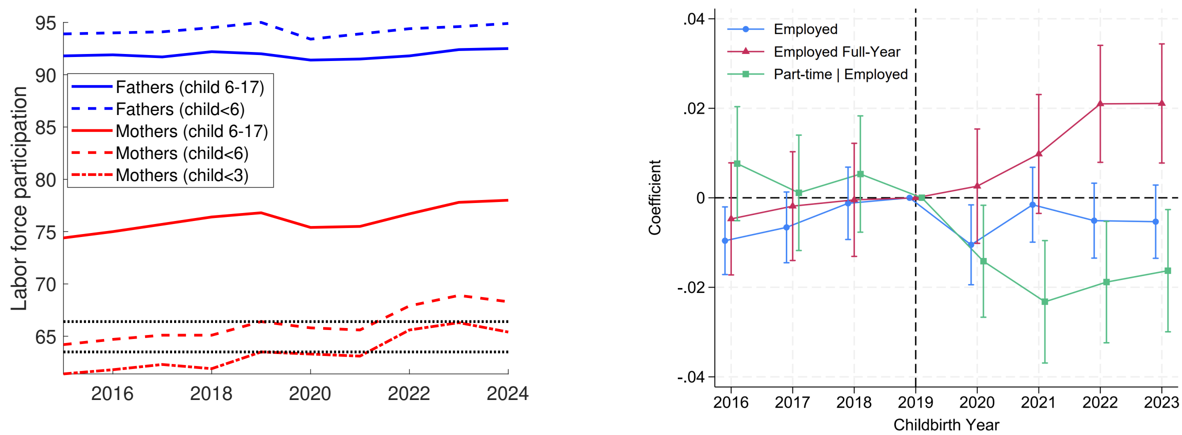

Note: Panel (a): CPS; change in labor force participation of mothers and fathers with young children relative to comparable adults without children. Panel (b): event-study estimates of the effect of remote-work feasibility on mothers' employment, from Basso et al. (2026).

5 Quantitative results

We analyze the impact of increased workplace flexibility—modeled as the measured reduction in the childcare cost parameter Φ—on labor supply, occupational choice, and gender gaps. We distinguish between the short-run transition dynamics, where existing cohorts are largely locked into previous occupational choices, and the long-run equilibrium, where new cohorts re-optimize their initial career paths and wages adjust.

5.1 Short- to medium-run transition dynamics

We first study the short- to medium-run impact of the WFH shock. We initialize the distribution of matches, occupational choices, and asset holdings at their baseline levels. In year one, the work-from-home parameter declines from 1 to 0.4, and we solve the household problem forward for five years holding prices fixed at their steady-state level. The frictional nature of occupational choice means that incumbents can adjust occupations only when a switching opportunity arrives, and new entrants have not yet re-optimized their initial choice.

Figure 12 presents the response of the share of women and men working in non-linear occupations over the life cycle, at one, three, and five years after the shock. Three patterns stand out. First, female employment and the female non-linear share rise on impact, but modestly: after five years female employment is up 1.7 percentage points and the female non-linear share 1.6 percentage points, against long-run responses several times larger. Second, the response is strongly concentrated among younger women (ages 25–35), who face the highest childcare burdens and have the longest horizon over which to reap the returns to human capital accumulation; the stock of older workers remains largely on pre-shock career paths. Third, men are essentially unaffected in the short run: male employment falls by 0.6 percentage points, hours are unchanged, and the male non-linear share declines by only 0.4 percentage points, driven by a small income effect in households where female earnings rise.

Figure 12: Short- to medium-run response: non-linear occupation shares over the life cycle

(a) Women (b) Men

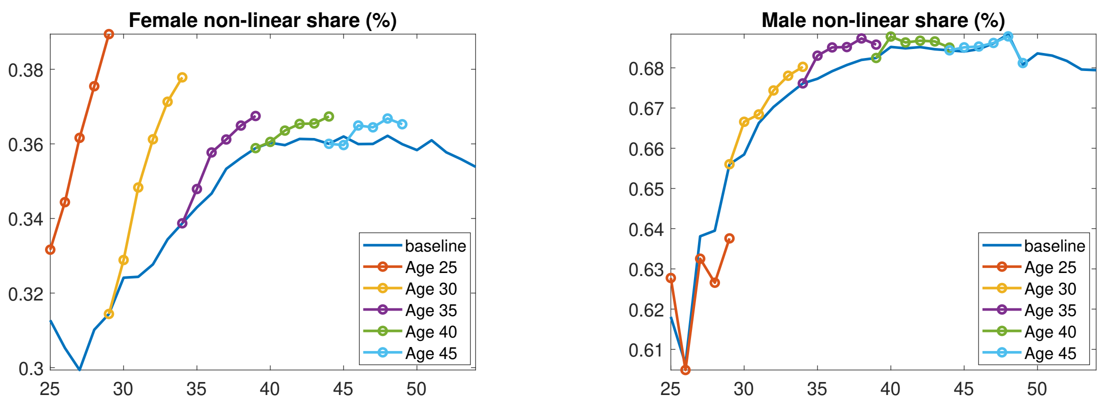

Note: Share of the population in non-linear occupations by age, in the baseline and at horizons of one to five years after the WFH shock. Prices fixed at baseline values; initial occupational choices of entrants held at baseline.

In the short run, most of the change in the occupational structure operates through the extensive labor supply margin rather than through switching. Women in linear occupations are the marginal workers: they increase participation by more than non-participating women attached to non-linear occupations, so that, conditional on employment, the female non-linear share actually falls slightly on impact even as the unconditional share rises. Female hours growth is correspondingly concentrated in the linear sector—exactly the pattern observed in the post-2020 data (Figure 7).

Validation against the post-pandemic data. The medium-run simulation provides an out-of-sample check of the model against the observed evolution of the labor market. Table 5 compares the change in female relative to male outcomes between 2018–19 and 2024–25 in the CPS, $\hat{X} = \Delta_{19\text{–}25} X_f - \Delta_{19\text{–}25} X_m$, with the model's five-year response. To respect the model's timing, we restrict the sample to workers aged 30 and above, who have already made their initial occupational choice. The model accounts for essentially all of the observed relative change in employment and around 60 percent of the relative shift toward non-linear occupations. It overshoots on the relative growth in full-time and overtime work. Notably, the model predicts only a small decline in relative female wages (-0.6 percent against -3.6 in the data), and the direction of the discrepancy is informative: both in the model and, plausibly, in the data, the entry of marginal female workers with below-average productivity depresses measured average wages—a selection effect that a composition-adjusted wage series would attenuate.

Table 5: Model validation: change in female relative to male outcomes, 2018–19 to 2024–25

	Data	Model	Share explained
Employment (p.p.)	2.18	2.30	1.06
Non-linear share (p.p.)	2.82	1.70	0.60
Full-time/overtime share (p.p.)	1.12	2.25	2.00
Wages (%)	−3.64	−0.64	0.18

Note: Data column reports $\hat{X} = \Delta_{19\text{–}25} X_f - \Delta_{19\text{–}25} X_m$ in the CPS for workers aged 30 and above, so that the initial occupation margin is closed. Model column reports the corresponding five-year partial-equilibrium response. Share explained is the ratio of the model to the data change.

5.2 Long-run occupational reallocation

In the long run, new cohorts enter the labor market anticipating the higher flexibility provided by WFH. This allows for a full adjustment of initial career choices, leading to substantial occupational reallocation and a much stronger overall response. In addition, the wages of linear and non-linear occupations are allowed to adjust in response to the changed supply of labor.

Initial sorting. Panel (a) of Figure 13 shows the joint occupational choices made by households before entering the labor market. In the baseline, the modal configuration pairs a man in a non-linear occupation with a woman in a linear occupation (35 percent of new households). In the WFH counterfactual this pattern changes markedly. There is a pronounced increase in women choosing non-linear occupations: the share of new households in which the woman holds the non-linear job and the man the linear job more than doubles, from 11 to 24 percent. Conversely, men shift towards linear occupations, and the share of traditional configurations falls from 35 to 19 percent. As childcare has effectively become less specialized, the household increasingly assigns the greedy occupation to the spouse with the higher expected return rather than by gender, and can select the lower-productivity member to supply flexibility. These new patterns of occupational sorting are further reinforced by the general equilibrium channel, which sees relative wages rise in the linear sector and pull men into flexible jobs.

Figure 13: Occupational sorting in the long run

(a) Initial joint occupational choice

(b) Life-cycle shares, conditional on employment

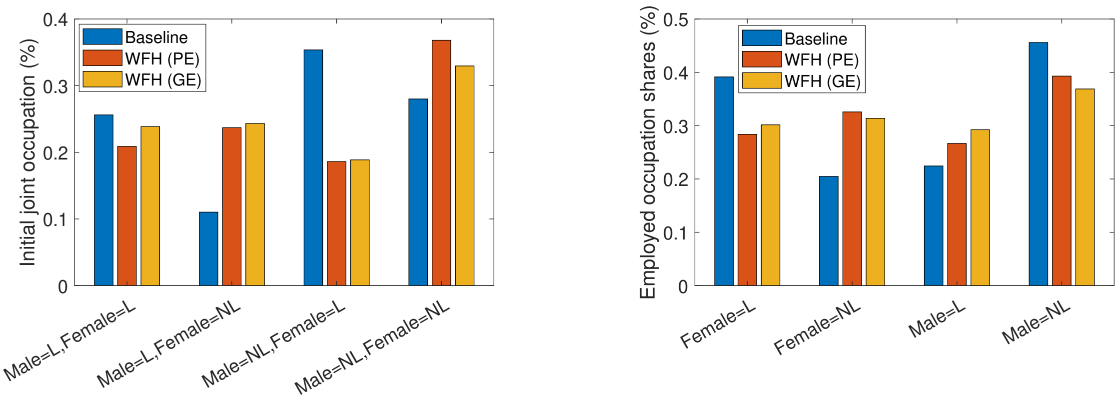

Note: Panel (a): joint occupational choices of newly formed households, baseline vs. WFH general equilibrium. Panel (b): occupational shares over the life cycle conditional on employment.

Life-cycle allocations. The initial pattern of sorting provides an incomplete picture, as members switch occupations along the life cycle and not all agents work. Panel (b) shows life-cycle occupational shares conditional on employment. The female non-linear share now starts higher at labor market entry and remains elevated throughout the life cycle—in contrast to the baseline, where women continuously exit high-intensity careers as family demands grow. Because men retain higher employment rates, they remain *more likely* than women to work in non-linear occupations over the life cycle, but the gap narrows sharply. Men initially reallocate away from non-linear occupations, though part of this is undone over the life cycle as some men switch back, reflecting both that men increasingly match into joint-linear households and that women retain a comparative advantage in childcare. Figure 14 decomposes the male reallocation by demographic group: the decline in the male non-linear share is concentrated among fathers of young children, while men in households with older children partially reallocate back toward non-linear occupations as childcare demands recede. Households thus time the supply of male flexibility to the period of peak childcare needs, rather than reducing male careers uniformly. In contrast to the short run, where the marginal female worker entered linear employment, in the long run most of the female adjustment flows from non-participation into non-linear employment.

Figure 14: Change in male non-linear share by demographic group

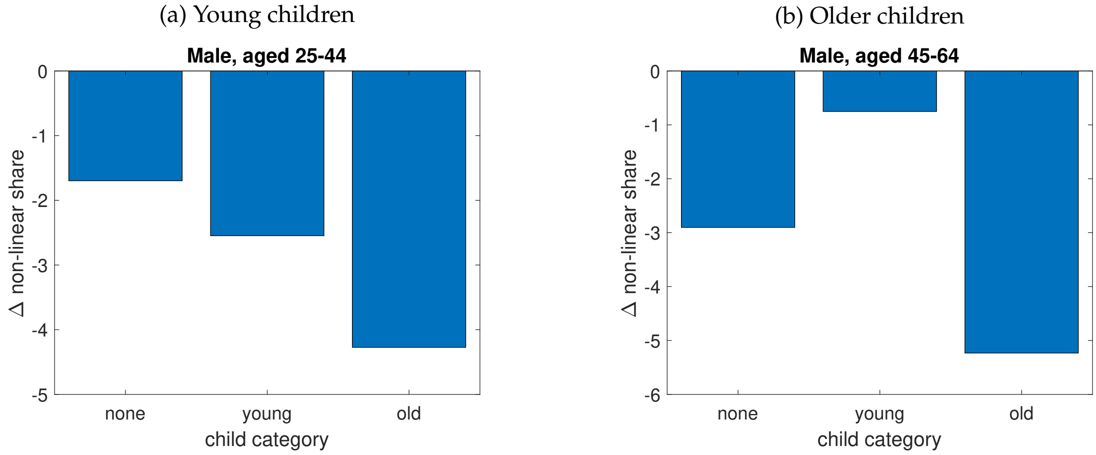

Note: Change in the male non-linear occupation share between the baseline and WFH steady states, by age and children's age.

Hours over the life cycle. Figure 15 displays the life-cycle implications. After households have fully adjusted, women supply 7.3 percent more hours (Table 6), partly in response to the higher returns to long hours in the non-linear sector. In the baseline, female hours exhibit a characteristic dip during the child-rearing years (ages 30–40); under the WFH regime this dip is substantially smoothed. Women maintain higher hours during peak childcare years, preventing the depreciation of human capital that typically exacerbates the child penalty. Offsetting this is a smaller decline in male labor supply (-3.4 percent) spread throughout the life cycle.

Figure 15: Long-run response over the life cycle

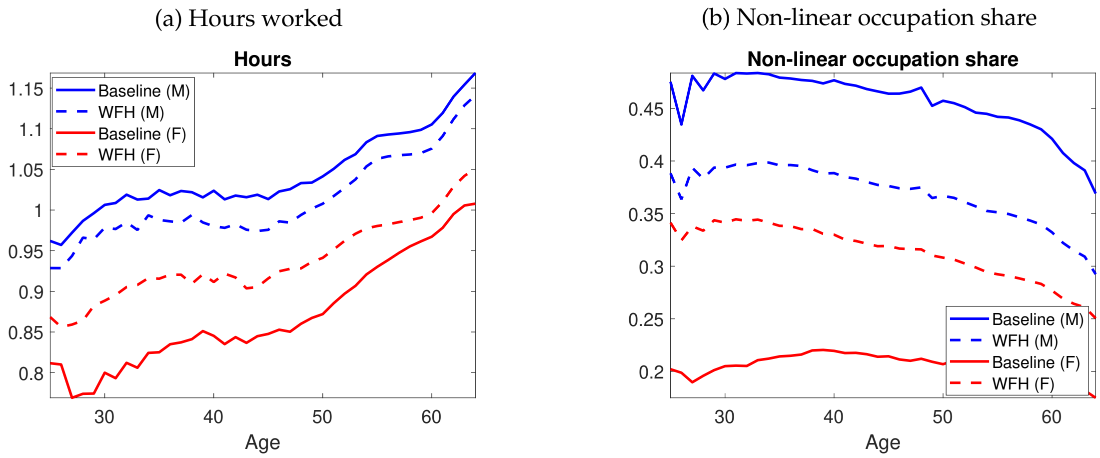

Note: Life-cycle profiles of hours worked and non-linear occupation shares by gender, baseline vs. WFH general equilibrium.

5.3 General equilibrium effects and the transition

The large-scale entry of women into non-linear occupations generates general equilibrium effects on wages. The increased supply of labor in the non-linear sector exerts downward pressure on the non-linear wage ω_{NL} and upward pressure on the linear wage ω_L: per efficiency unit, the non-linear wage falls by 0.7 percent and the linear wage rises by 0.5 percent, a 1.2 percent shift in relative wages.

Table 6 presents aggregate responses at the five-year horizon and in the long run, with and without general equilibrium price adjustment. Three sets of results stand out.

Table 6: Aggregate responses to the work-from-home shock

Change	Male			Female		
	5 years	P.E.	G.E.	5 years	P.E.	G.E.
Income (%)	−0.6	−4.9	−6.9	4.3	13.5	12.8
Employment (p.p.)	−0.6	−2.1	−1.9	1.7	1.4	2.0
Employment with young child (p.p.)	−3.7	−4.4	−4.3	9.6	7.8	8.5
Hours (%)	−0.2	−2.3	−3.4	2.6	8.2	7.3
Wages (%)	0.4	0.4	−0.5	−0.2	3.6	3.1
Non-linear share (p.p.)	−0.4	−6.3	−8.7	1.6	12.1	10.9

Note: "5 years" is the medium-run partial-equilibrium transition of Section 5.1. "P.E." is the long-run steady state holding wages at baseline values; "G.E." allows wages to adjust. Changes relative to the baseline steady state; p.p. denotes percentage points.

First, the headline finding: in the long run, male income falls by 6.9 percent while female income rises by 12.8 percent, so that total household income rises modestly. The long-run female income gain is three times the medium-run response after five years (4.3 percent), and the occupational reallocation is roughly seven times larger (10.9 versus 1.6 percentage points), underscoring the importance of cohort turnover and the entry margin. Income rises through a combination of employment, hours, and wages, with reallocation toward the non-linear sector the dominant contributor for women. The employment rate of women with young children rises by 8.5 percentage points in the long run, closing around two-thirds of the baseline motherhood employment gap, while the employment of fathers with young children falls by 4.3 percentage points—a within-household substitution that is essentially absent at the five-year horizon.

Second, general equilibrium price adjustments matter, but in a direction-specific way. For women, the falling non-linear wage dampens reallocation: the female non-linear share rises 10.9 percentage points in general equilibrium against 12.1 in partial equilibrium, and female income gains are trimmed from 13.5 to 12.8 percent. For men, by contrast, general equilibrium *amplifies* the exit from non-linear occupations, from −6.3 to −8.7 percentage points: the compressed non-linear premium and the rising linear wage reinforce the household-level incentive for men to supply flexibility. Because women remain over-represented in linear occupations, the rise in the linear wage also directly narrows gender wage gaps. Ignoring the general equilibrium channel would thus overstate female reallocation, understate male reallocation, and mischaracterize the distribution of the gains within households.

Third, the transition is protracted. Figure 16 plots the transition paths of consumption, hours, and the non-linear employment share following the shock; the computation of the equilibrium price path is described in Appendix C.3. Aggregates converge to the new steady state only over decades, as successive entering cohorts embody the new occupational choices and accumulate human capital along their new career paths. While the extensive employment margin adjusts quickly, the income effects at five years are a third of their eventual magnitude and the occupational reallocation less than a fifth. Evaluations of flexible work based on short post-pandemic windows will therefore capture only a small fraction of its eventual impact.

Figure 16: Transition path to the new work-from-home steady state

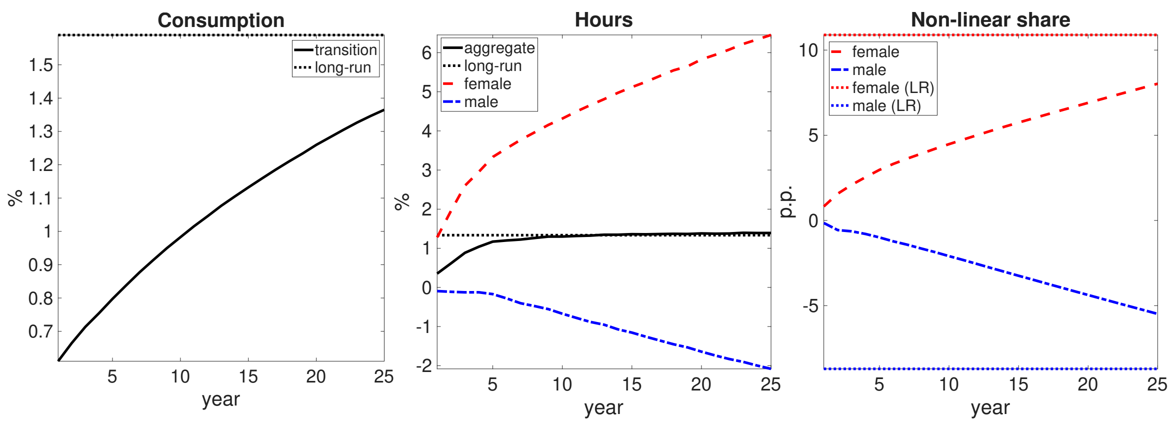

Note: Transition paths of aggregate consumption, hours by gender, and the non-linear employment share by gender following the permanent WFH shock.

Appendix Table A.1 reports an alternative experiment in which the WFH gain is scaled by each occupation's teleworkability, so that the flexibility gain is approximately 30 percent larger in the non-linear sector. The asymmetric shock generates a larger shift of female labor supply toward non-linear occupations and correspondingly larger general equilibrium wage responses, but the aggregate picture is very similar.

6 Welfare and the career cost of children

We now use the model to quantify the welfare consequences of the rise in workplace flexibility. Table 7 collects the results, all measured in the long-run general equilibrium. We first define the two objects the table reports.

Measurement. Our welfare metric is the expected value of a newly formed household at labor-market entry. For economy $e \in \{\text{Baseline}, \text{WFH}\}$, define

$$\bar{V}^{e} \;=\; \sum_{o_f, o_m} \Pr^{e}(o_f, o_m)\; \hat{V}^{e, o_f, o_m}, \tag{16}$$

where $\hat{V}^{e, o_f, o_m}$ is the ex-ante value of a household with occupational composition (o_f, o_m) defined in equation (11), evaluated at the prices and decision rules of economy e, and $\Pr^{e}(o_f, o_m)$ is the equilibrium joint distribution of initial occupational choices in that economy. The first row of Table 7 reports the proportional change $(\bar{V}^{\text{WFH}} - \bar{V}^{\text{Base}})/|\bar{V}^{\text{Base}}|$. Because flow utility is negative under our CRRA specification, this change has no direct cardinal interpretation; we therefore convert it into an asset-equivalent variation, Δ^{a}, defined as the proportional increase in the initial asset endowment that would deliver the same expected value within the baseline economy:

$$\bar{V}^{\text{Base}}\big((1+\Delta^{a})\,\bar{a}_0\big) \;=\; \bar{V}^{\text{WFH}}, \tag{17}$$

where $\bar{V}^{\text{Base}}(a_0)$ traces the baseline expected value as a function of the initial asset endowment and $\bar{a}_0$ denotes its baseline level. A natural alternative would be the standard consumption-equivalent variation. In our setting, however, the warm-glow bequest motive $B(a')$ makes preferences non-homothetic: a uniform proportional increase in consumption in all dates and states does not scale the value function, so the consumption-equivalent has no closed form and would require re-solving the household problem for each candidate scaling, conflating the valuation of consumption with the revaluation of bequests. The asset-equivalent variation is well-defined in this non-homothetic environment and answers a concrete question: the lump-sum wealth transfer at labor-market entry that would make a baseline household as well off as the arrival of flexible work. Because newly formed households hold few assets relative to their lifetime resources, the percentage figures are large; expressed in flow terms, the transfer corresponds to approximately 21 percent of average annual household income. Gender-specific gains are computed analogously from the ex-ante values of men and women prior to matching, inclusive of their occupational choice margins.

The career cost of children compares each economy to a counterfactual version of *itself* in which children are shut down: the probability of entering the labor market with a child and all fertility hazards are set to zero, household decision rules are re-solved, and the economy is simulated holding prices and all other primitives at the values of economy e. For an outcome x, define its net present value at labor-market entry,

$$\mathcal{N}^{e}(x) \;=\; \sum_{j=1}^{J} \beta^{j-1}\, m_j\, \bar{x}^{e}_{j}, \tag{18}$$

where $\bar{x}_j^e$ is the cross-sectional mean of x at age j in economy e and m_j is the surviving share of the cohort at age j. The career cost of children in outcome x is the share of this value lost to the presence of children,

$$\mathcal{CC}^e(x) \;=\; 1 - \frac{\mathcal{N}^e(x)}{\mathcal{N}^e_{\varnothing}(x)}, \tag{19}$$

where the subscript $\varnothing$ denotes the no-children counterfactual. We compute $\mathcal{CC}$ for household consumption per equivalence-scale unit over the full life cycle ($J = J^W + J^R$) and for the labor income of each spouse over the working life ($J = J^W$). Panel B of Table 7 reports the change in these costs induced by the WFH shock, $\mathcal{CC}^{\text{WFH}}(x)/\mathcal{CC}^{\text{Base}}(x) - 1$.

Table 7: Welfare and the career cost of children

	Change (%)
A. Welfare	
Expected value of household	1.8
In terms of initial average asset value	203.9
B. Career cost of children	
NPV consumption	-11.7
NPV male income	35.9
NPV female income	-19.7
C. Misallocation and risk	
Female misallocation (initial z draw)	-57.6
Full-time female wage penalty	-44.4
Consumption growth variance	-5.3
Hours growth variance	-16.3

Note: Long-run general equilibrium changes relative to baseline. Panel A: change in the expected value of a newly formed household, equation (16), and its asset-equivalent variation, equation (17). Panel B: change in the career cost of children, $\mathcal{CC}(x)$ of equation (19). Panel C: female misallocation is the incidence of households in which the wife receives the higher initial productivity draw but does not hold the non-linear occupation; the full-time female wage penalty is the average within-occupation wage gap for full-time workers; variances are of household consumption growth and individual hours growth in the simulated panel.

Welfare. The expected value of a newly formed household, equation (16), rises by 1.8 percent. The asset-equivalent variation of equation (17) expresses this in interpretable units: flexibility is worth approximately a 204 percent increase in the average initial asset endowment, even after general equilibrium wage adjustments (the partial equilibrium figure is somewhat larger, as wage compression in the non-linear sector redistributes part of the gain). Gains are similar for men and women: although male earnings fall, household consumption rises and the burden of supplying long inflexible hours is shared more efficiently. This near-equality partly reflects the unitary framework, in which resources are pooled and the household fully internalizes the redistribution of the career cost of children toward fathers. In collective or limited-commitment settings (Mazzocco, 2007;

Voena, 2015; Lise and Yamada, 2019), the decline in male relative earnings and the rise in the male career cost of children would shift bargaining power within the household—plausibly reinforcing the reallocation we document, as women's outside options improve, but opening individual-level welfare differences and renegotiation or divorce margins on which the unitary model is silent. Our gender-specific ex-ante values, computed prior to matching and inclusive of the occupational choice margin, partially address this concern by measuring gains at the stage at which interests are not yet pooled; a full treatment of intra-household bargaining over flexibility is an important avenue for future work.

The career cost of children. Applying equation (19) in the baseline economy, children cost the household 14.1 percent of the net present value of lifetime consumption, $\mathcal{CC}^{\text{Base}}(c) = 0.141$; the cost is starkly gendered, amounting to 40.1 percent of the mother's lifetime income against 13.9 percent of the father's. The WFH shock reduces the consumption cost of children by 11.7 percent and the female-income cost by 19.7 percent (from 40.1 to 32.2 percent of lifetime income). The male-income cost *rises* by 35.9 percent: within the household, fathers absorb a larger share of the career cost as couples reallocate childcare and hours—a redistribution, not a loss, from the household's perspective.

Misallocation and risk. The welfare gain reflects more than higher female earnings. First, flexibility improves the allocation of talent within households: the incidence of misallocation—households in which the wife draws the higher initial productivity but the husband occupies the non-linear job—falls by 58 percent. Second, the full-time female wage penalty falls by around 45 percent, as women maintain human capital through the child-rearing years. Third, flexibility acts as insurance: households can smooth labor supply around child-related shocks rather than adjusting through costly consumption movements, so the variance of consumption growth falls by 5.3 percent and the variance of hours growth by 16.3 percent.

Decomposing the female-income cost of children. To understand which margins drive the reduction in the career cost of children, we decompose female income at each age into the product of four components:

$$Y_j = \sum_o \underbrace{\Pr(o)\,\omega_o\,\ell(o,j)}_{\text{occupation}} \times \underbrace{\bar{Z}_o}_{\text{productivity}} \times \sum_h \underbrace{h^{\theta_o}\Pr(h|h>0,o)}_{\text{hours}} \times \underbrace{\Pr(h>0|o)}_{\text{employment}} \tag{20}$$

and attribute the gap $\mathcal{N}_{\varnothing}(y_f) - \mathcal{N}(y_f)$ of equation (18) to each component. Table 8 reports the shares before and after the WFH shock.

Two findings emerge. First, the importance of the extensive margin declines (from 42.6 to 32.9 percent): flexible work keeps mothers employed, so non-participation accounts for a smaller share of the remaining penalty. Second, the residual cost increasingly reflects dynamics: the productivity component—slower human capital accumulation of mothers—rises to become the dominant contributor (83.0 percent), and rises relative to the occupation

Table 8: Decomposition of the female-income career cost of children

	Work from home	
Share (%)	Pre	Post
Occupation	22.5	32.7
Productivity	53.9	83.0
Hours	−30.6	−53.2
Employment	42.6	32.9
Interaction	11.6	4.6

Note: Contribution of each component of equation (20) to the gap in the net present value of female income between households without and with children, before and after the WFH shock. The interaction row collects the cross terms of the multiplicative decomposition; shares sum to 100.

component. The negative hours contribution indicates that, conditional on working and on occupation, mothers' hours choices per se would imply a smaller income gap; the hours cost of children operates through occupational composition and its dynamic consequences rather than through hours conditional on occupation. In short, as flexibility removes the participation barrier, what remains of the child penalty is a human capital phenomenon. The interaction term is modest and shrinks after the shock (from 11.6 to 4.6 percent), indicating that the four margins capture the gap without large offsetting cross effects.

7 Artificial intelligence and the future of flexible work

Work-from-home is unlikely to be the last technology to reshape the occupational constraints facing households. Artificial intelligence is widely expected to shift labor demand across occupations over the coming decade, and early evidence from high-frequency payroll data indicates that this shift has already begun: Brynjolfsson, Chandar, and Chen (2025) document substantial relative employment declines for early-career workers in the occupations most exposed to generative AI, concentrated in uses of AI that automate rather than augment work—precisely the labor-market entry margin through which long-run adjustment operates in our model. The incidence of AI exposure is, moreover, far from gender-neutral. In this section we first document that AI exposure is systematically related to the greedy-job and teleworkability classifications at the heart of our analysis, and then use the model's task-based production structure to quantify how an AI-driven demand shift interacts with the gains from flexible work.

7.1 AI exposure across occupations and genders

A striking fact is that the occupations that benefit most from teleworking are also the most exposed to AI. Panel (a) of Figure 17 plots occupation-level AI exposure, measured by the Felten, Raj, and Seamans (2021) AI Occupational Exposure index, against the Dingel and Neiman (2020) teleworkability measure: the correlation is strongly positive. Appendix

Figure A.4 shows the relationship is robust across four alternative exposure measures, including the language-model task exposure of Eloundou et al. (2024), machine-learning suitability (Brynjolfsson, Mitchell, and Rock, 2018), and AI patent exposure (Webb, 2020). The technology that relaxed the time constraints of greedy jobs may thus simultaneously erode their labor demand.

The incidence of AI exposure is also sharply gendered. Panel (b) of Figure 17 plots occupation-level exposure from Eloundou et al. (2024) against the occupation's hours rank, separately for the occupations employing men and women. Male and female *non-linear* occupations are similarly exposed. The difference lies in the linear sector: the AI exposure of the linear occupations in which women are concentrated—administrative, clerical, and related work—is almost double that of male linear occupations, which are weighted toward manual and physical tasks. Appendix Figures A.5 and A.6 show the pattern is robust to using the AIOE measure, to weighting by earnings, and is visible across the entire exposure distribution.

Figure 17: AI exposure, teleworkability, and gender

(a) Teleworkability and AI exposure

(b) AI exposure by occupation hours rank and gender

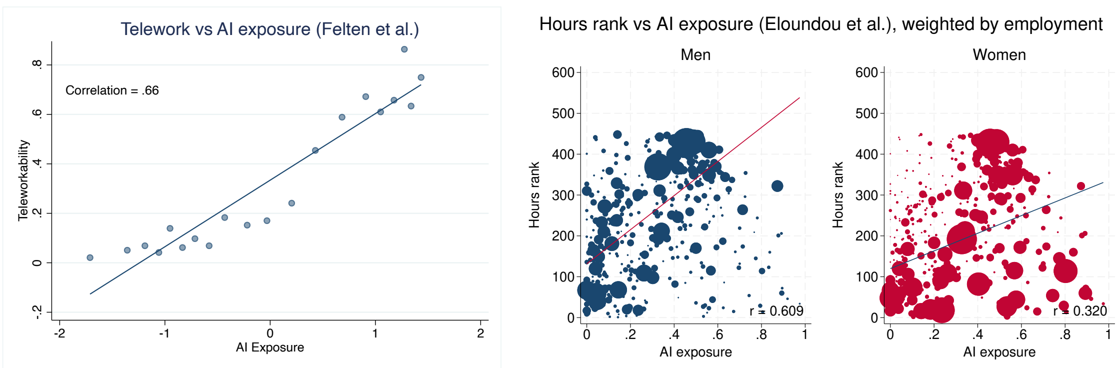

Note: Panel (a): AI occupational exposure defined as in Felten, Raj, and Seamans (2021); teleworkability as in Dingel and Neiman (2020). Panel (b): occupation-level AI exposure from Eloundou et al. (2024) against the occupation's male-hours rank, weighted by employment, separately for men and women.

7.2 The AI experiment

We feed the measured task-level AI exposure into the production technology of Section 3.4. Exposure by task, occupation, and gender from Eloundou et al. (2024) is aggregated to the model's occupation-by-gender cells and mapped into the share of each cell's tasks that is automated; in the baseline experiment, we assume that half of exposed tasks are ultimately automated (Appendix B details the mapping). Automation reallocates tasks from labor to capital, reducing the effective labor demand of exposed cells, while the associated cost saving raises aggregate productivity (Acemoglu, 2025). We compute two AI counterfactuals: an economy in which AI arrives on top of the work-from-home equilibrium (WFH+AI), and, as a reference, an economy in which AI arrives in the pre-pandemic baseline without flexible work (AI only). Comparing WFH+AI to WFH isolates the marginal effect of AI in a flexible-work economy, while the AI-only economy reveals how much of that effect depends on the presence of flexibility—the interaction between the two shocks that is the focus of this section. We deliberately abstract from any effect of AI on the flexibility parameter itself.

Table 9 reports the results. The AI demand shift lowers incomes for both genders by around 4 percent, reflecting the direct displacement effect net of the productivity gain. But its allocative effects are gendered and operate exactly through the mechanism at the core of the model: because female linear occupations are the most exposed, AI erodes the return to precisely the flexible jobs that women have traditionally used to reconcile work and family. The result is a further reallocation of women toward non-linear occupations (the female non-linear share rises by an additional 0.5 percentage point, while the male share falls by 2.2) and a further 16.5 percent decline in the full-time female wage penalty. At the same time, the degraded value of part-time linear work removes an important margin of adjustment for mothers: aggregate female employment is essentially unchanged, but the motherhood employment penalty widens by 0.1 percentage point relative to the WFH economy.

Table 9: Effect of AI, relative to the work-from-home equilibrium

	Change, WFH $\rightarrow$ WFH+AI	
	Male	Female
Income (%)	-4.0	-3.7
Non-linear share (p.p.)	-2.2	0.5
Employment (p.p.)	0.5	0.1
Hours (%)	-0.4	0.2
Full-time female wage penalty (%)		-16.5
Motherhood employment penalty (p.p.)		0.1

Note: Long-run general equilibrium changes of the WFH+AI economy relative to the WFH economy. A positive motherhood employment penalty entry denotes a widening of the penalty.

7.3 The interaction of AI and flexibility

Table 10 summarizes the four steady states: baseline, AI only, WFH, and WFH+AI. The central lesson is an interaction effect. On its own, AI barely moves the family margins at the core of this paper: in the AI-only economy the motherhood employment gap is essentially unchanged relative to baseline (a change of less than 0.1 percentage point), and the female non-linear share rises only modestly, because without flexible work the childcare constraint continues to keep mothers out of greedy occupations regardless of relative wages. It is the combination of the two technologies that is potent. With flexibility in place, AI erodes the return to precisely the female-dominated linear occupations that mothers have traditionally used to reconcile work and family, pushing women further into non-linear work (the female non-linear share rises to 51.7 percent, from 51.0 under WFH alone) while degrading the value of the part-time linear option: the employment gap of mothers with young children, which WFH had narrowed from -15.4 to -5.8 percentage points, widens again to -6.5, and the female-income cost of children rises from 32.2 back to 34.8 percent—reversing roughly a third of the WFH gain.

Table 10: Summary: baseline, AI, work-from-home, and combined steady states

	Baseline	AI only	WFH	WFH+AI
Cost of children, consumption (%)	14.1	–	12.4	13.9
Cost of children, female income (%)	40.1	–	32.2	34.8
Non-linear share of women (%)	34.3	37.2	51.0	51.7
Employment gap, women with young children (p.p.)	-15.4	-15.3	-5.8	-6.5
Expected household value relative to baseline (%)	–	-2.2	1.8	-0.3
Value of flexibility within each economy (%)	1.8		2.1	

Note: Long-run general equilibrium values. The non-linear share of women is conditional on employment. The employment gap is for women with young children relative to comparable women without children. The penultimate row reports the change in the expected value of a newly formed household, equation (16), relative to the pre-pandemic baseline. The final row reports the value of flexibility—the gain from WFH holding the AI regime fixed—in the pre-AI economy (baseline $\rightarrow$ WFH) and in the AI economy (AI $\rightarrow$ WFH+AI). The cost of children in the AI-only economy is not computed.

The welfare accounting makes the same point in a different way. Relative to the pre-pandemic baseline, the combined WFH+AI economy delivers an expected household value that is approximately unchanged (-0.3 percent, or -27 percent in asset-equivalent terms). But this aggregate figure is the sum of two conceptually distinct components: a negative level effect of the AI demand shift (-2.2 percent of expected household value in the AI-only economy) and the value of flexibility, which is essentially *undiminished* by AI—comparing the AI economy with and without WFH yields a flexibility gain of 2.1 percent of expected household value, if anything slightly above the 1.8 percent in the pre-AI economy. AI is a level shock operating through labor demand, not a reversal of the flexibility channel. We caution against reading the level effect as a forecast: it scales directly with the assumed automation share (one half of exposed tasks) and abstracts from AI-driven productivity gains beyond the Hulten term, so its magnitude is illustrative. The robust and, we believe, economically meaningful findings are the interactions: flexibility is what makes AI matter for the family margins—opening greedy jobs to women while AI erodes the female-dominated linear alternative—and AI in turn reshapes where the gains from flexibility accrue, reinforcing female sorting into high-return occupations even as it claws back part of the reduction in the career cost of children. The family-related gains from remote work therefore depend in part on the evolution of labor demand in the occupations that flexibility made accessible.

8 Conclusion

We develop a dynamic general equilibrium life-cycle model of household labor supply and occupational choice to quantify the role of workplace flexibility in determining the gender earnings gap. By modeling the trade-off between "greedy" occupations—which reward long hours but lack flexibility—and flexible "linear" occupations, we replicate the persistent sorting of women into lower-paying jobs and the resulting stagnation of female human capital. New empirical evidence disciplines the analysis: greedy occupations are disproportionately teleworkable, post-pandemic female employment growth is concentrated in teleworkable greedy occupations, and occupational switching over the working life is rare.

Our counterfactual analysis of the post-pandemic rise in work-from-home yields four key insights. First, labor market adjustments are sluggish: in the short run, existing cohorts are locked into prior career choices, and the five-year response—which we validate against the observed 2019–2025 evolution of relative female outcomes—is a fraction of the eventual impact. Second, in the long run, the reduction in effective childcare costs drives a substantial reallocation: new cohorts of women increasingly select into high-return, non-linear occupations, female earnings rise by 12.8 percent, and the career cost of children in female earnings falls by a fifth. Third, general equilibrium effects are crucial: the influx of women into greedy jobs compresses the non-linear premium, which dampens female reallocation but amplifies the reallocation of men toward flexible jobs and reinforces the closing of gender income gaps. Fourth, the welfare gains are large—equivalent to roughly a 200 percent increase in initial assets—and derive not only from higher earnings but from a better allocation of talent within households and reduced exposure to consumption and hours risk.

Looking forward, we show that the gains from flexibility are exposed to the next wave of technological change. The occupations that benefit from teleworking are the most exposed to artificial intelligence, and female-dominated flexible occupations are the most exposed of all. Feeding measured AI exposure into the model, we find that AI and flexibility interact strongly: on its own, AI leaves the family margins essentially untouched, but in combination with flexible work it reverses roughly a third of the WFH gains in the career cost of children even as it reinforces female sorting into high-return occupations—while leaving the welfare value of flexibility itself undiminished. These findings suggest that while workplace reorganization toward greater flexibility can reduce persistent gender disparities, the full benefits will materialize slowly as new cohorts replace older ones, and their durability depends on how labor demand evolves. Future work should consider policy design in this environment—including joint taxation and childcare provision—as well as richer evidence from administrative data on the household-level reallocation we highlight.

References

Aaronson, Daniel and Eric French. 2004. "The effect of part-time work on wages: Evidence from the social security rules." *Journal of Labor Economics* 22 (2):329–352.

Acemoglu, Daron. 2025. "The Simple Macroeconomics of AI." *Economic Policy* 40 (121):13–58.

Acemoglu, Daron and Pascual Restrepo. 2018. "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review* 108 (6):1488–1542.

Adams-Prassl, Abi, Teodora Boneva, Marta Golin, and Christopher Rauh. 2022. "Work that can be done from home: Evidence on variation within and across occupations and industries." *Labour Economics* 74:102083.

Adda, Jérôme, Christian Dustmann, and Katrien Stevens. 2017. "The career costs of children." *Journal of Political Economy* 125 (2):293–337.

Aksoy, Cevat Giray, Jose Maria Barrero, Nicholas Bloom, Steven J. Davis, Mathias Dolls, and Pablo Zarate. 2022. "Working from Home Around the World." *Brookings Papers on Economic Activity* :281–360.

Albanesi, Stefania and Claudia Olivetti. 2007. "Gender Roles and Technological Progress." NBER Working Paper 13179, National Bureau of Economic Research.

———. 2009. "Home Production, Market Production and the Gender Wage Gap: Incentives and Expectations." *Review of Economic Dynamics* 12 (1):80–107.

Alon, Titan, Matthias Doepke, Jane Olmstead-Rumsey, and MichÃ¨le Tertilt. 2020. "The Impact of COVID-19 on Gender Equality." NBER Working Papers 26947, National Bureau of Economic Research, Inc. URL https://ideas.repec.org/p/nbr/nberwo/26947.html.

Ashraf, Nava, Oriana Bandiera, Virginia Minni, and Víctor Quintas-Martínez. 2023. "Gender and the Misallocation of Labor Across Countries." Working paper.

Atkin, David, Antoinette Schoar, and Sumit Shinde. 2023. "Working from Home, Worker Sorting and Development." Working Paper 31515, National Bureau of Economic Research.

Attanasio, Orazio, Hamish Low, and Virginia Sánchez-Marcos. 2008. "Explaining Changes in Female Labor Supply in a Life-Cycle Model." *American Economic Review* 98 (4):1517–1552.

Autor, David H, Lawrence F Katz, and Melissa S Kearney. 2008. "Trends in US wage inequality: Revising the revisionists." *The Review of economics and statistics* 90 (2):300–323.

Bang, Minji. 2022. "Job flexibility and household labor supply: Understanding gender gaps and the child wage penalty." Tech. rep., University of Pennsylvania.

Barrero, Jose Maria, Nicholas Bloom, and Steven J Davis. 2021. "Why working from home will stick." Tech. rep., National Bureau of Economic Research.

Basso, Gaetano, Maria De Paola, Salvatore Lattanzio, and Matteo Paradisi. 2026. "Workplace Flexibility and the Motherhood Penalty: Evidence from the Diffusion of Remote Work." Discussion Paper 21486, CEPR.

Berresheim, Ursula. 2025. "Work from Home, Work for Less? How workplace flexibility affects mothers careers." *working paper* .

Bick, Alexander, Adam Blandin, and Karel Mertens. 2022. "Work from home after COVID: The case of US micro data." *American Economic Journal: Macroeconomics* 14 (3):198–232.

Blau, Francine D. and Lawrence M. Kahn. 2017. "The Gender Wage Gap: Extent, Trends, and Explanations." *Journal of Economic Literature* 55 (3):789–865.

Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten. 2018. "Children, time allocation, and consumption insurance." *Journal of Political Economy* 126 (S1):S73–S115.

Braun, Christine and Ana Figueiredo. 2026. "Labor Market Beliefs and the Gender Wage Gap." Working papers. URL https://drive.google.com/file/d/15_9s3yg_zx0bZalBgWpzoQVUf2GgEfbT/view.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. 2025. "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence." Working paper, Stanford Digital Economy Lab.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock. 2018. "What Can Machines Learn and What Does It Mean for Occupations and the Economy?" *AEA Papers and Proceedings* 108:43–47.

Carter, Braxton, Kyle Herkenhoff, Jonathan Rothbaum, and Lawrence Schmidt. 2025. "Changing Income Risk Across the U.S. Skill Distribution: Evidence from a Generalized Kalman Filter." *American Economic Review* 115 (12):4438–4475.

Cortés, Patricia and Jessica Pan. 2023. "Children and the Remaining Gender Gaps in the Labor Market." *Journal of Economic Literature* 61 (4):1359–1409.

Dingel, Jonathan I and Brent Neiman. 2020. "How many jobs can be done at home?" *Journal of Public Economics* 186:104130.

Doepke, Matthias and Michèle Tertilt. 2016. "Families in Macroeconomics." In *Handbook of Macroeconomics*, vol. 2. Elsevier, 1789–1891.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2024. "GPTs are GPTs: Labor Market Impact Potential of LLMs." *Science* 384 (6702):1306–1308.

Emanuel, Natalia, Emma Harrington, and Amanda Pallais. 2026. "The Power of Proximity to Coworkers." *Quarterly Journal of Economics* Forthcoming.

Erosa, Andrés, Luisa Fuster, Gueorgui Kambourov, and Richard Rogerson. 2022. "Hours, occupations, and gender differences in labor market outcomes." *American Economic Journal: Macroeconomics* 14 (3):543–590.

Faberman, Jason, Andreas I. Mueller, and Aysegul Sahin. 2026. "Job Search, Job Amenities and the Gender Pay Gap." NBER Working Papers 34877, National Bureau of Economic Research, Inc. URL https://ideas.repec.org/p/nbr/nberwo/34877.html.

Felten, Edward, Manav Raj, and Robert Seamans. 2021. "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses." *Strategic Management Journal* 42 (12):2195–2217.

Gibbs, Michael, Friederike Mengel, and Christoph Siemroth. 2023. "Work from Home and Productivity: Evidence from Personnel and Analytics Data on Information Technology Professionals." *Journal of Political Economy Microeconomics* 1 (1):7–41.

Goldin, Claudia. 2014. "A Grand Gender Convergence: Its Last Chapter." *American Economic Review* :1091–1119.

Goldin, Claudia, Sari Pekkala Kerr, and Claudia Olivetti. 2024. "The history of women's work and wages." *Journal of Economic Dynamics and Control* :104593.

Greenwood, Jeremy, Nezih Guner, Georgi Kocharkov, and Cezar Santos. 2016. "Technology and the Changing Family: A Unified Model of Marriage, Divorce, Educational Attainment, and Married Female Labor-Force Participation." *American Economic Journal: Macroeconomics* 8 (1):1–41.

Harrington, Emma and Matthew E Kahn. 2025. "Has the Rise of Work from Home Reduced the Motherhood Penalty in the Labor Market?" Tech. rep., National Bureau of Economic Research.

Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow. 2019. "The Allocation of Talent and U.S. Economic Growth." *Econometrica* 87 (5):1439–1474.

Huckfeldt, Christopher. 2022. "Understanding the Scarring Effect of Recessions: The Role of Fluctuations in Job Quality." *American Economic Review* 112 (4):1273–1310.

Kleven, Henrik, Camille Landais, and Gabriel Leite-Mariante. 2025. "The Child Penalty Atlas." *The Review of Economic Studies* 92 (5):3174–3207.

Kleven, Henrik, Camille Landais, and Jakob Egholt Søgaard. 2019. "Children and Gender Inequality: Evidence from Denmark." *American Economic Journal: Applied Economics* 11 (4):181–209.

Krusell, Per, Lee E. Ohanian, José-Víctor Ríos-Rull, and Giovanni L. Violante. 2000. "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis." *Econometrica* 68 (5):1029–1053.

Lise, Jeremy and Ken Yamada. 2019. "Household Sharing and Commitment: Evidence from Panel Data on Individual Expenditures and Time Use." *Review of Economic Studies* 86 (5):2184–2219.

Mazzocco, Maurizio. 2007. "Household Intertemporal Behaviour: A Collective Characterization and a Test of Commitment." *Review of Economic Studies* 74 (3):857–895.

Traiberman, Sharon. 2019. "Occupations and Import Competition: Evidence from Denmark." *American Economic Review* 109 (12):4260–4301. URL https://www.aeaweb.org/articles?id=10.1257/aer.20161925.

Voena, Alessandra. 2015. "Yours, Mine, and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples?" *American Economic Review* 105 (8):2295–2332.

Webb, Michael. 2020. "The Impact of Artificial Intelligence on the Labor Market." Working paper, Stanford University.

Wu, Chunzan and Dirk Krueger. 2021. "Consumption insurance against wage risk: Family labor supply and optimal progressive income taxation." *American Economic Journal: Macroeconomics* 13 (1):79–113.

Appendix

A Additional empirical evidence

Figure A.1: Teleworkability and occupation ranks: robustness

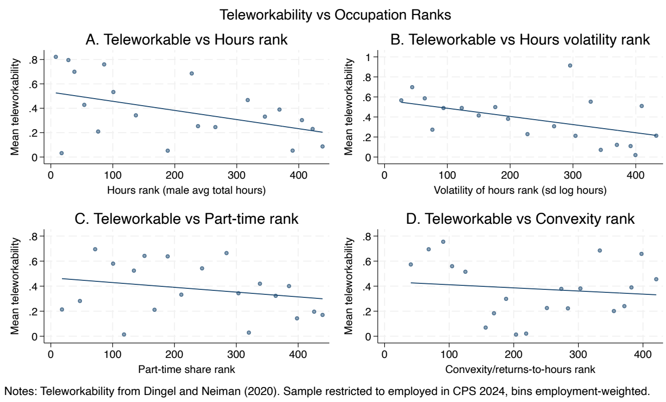

Note: Occupations classified as teleworkable following Dingel and Neiman (2020), based on the feasibility of performing tasks remotely using O*NET data.

Figure A.2: Employment shares in teleworkable occupations by year and gender

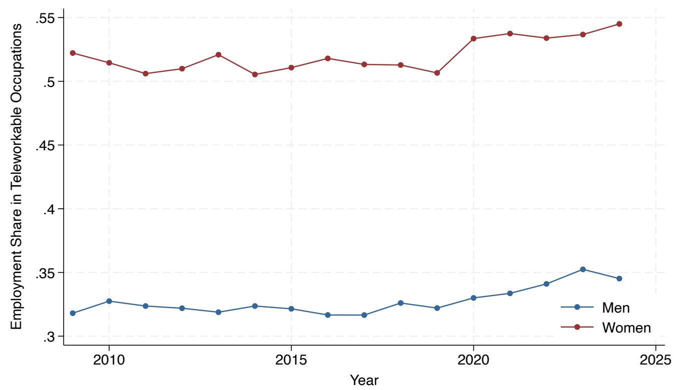

Note: Occupations classified as teleworkable following Dingel and Neiman (2020). CPS.

Figure A.3: Joint occupational choice at the household level around 2020

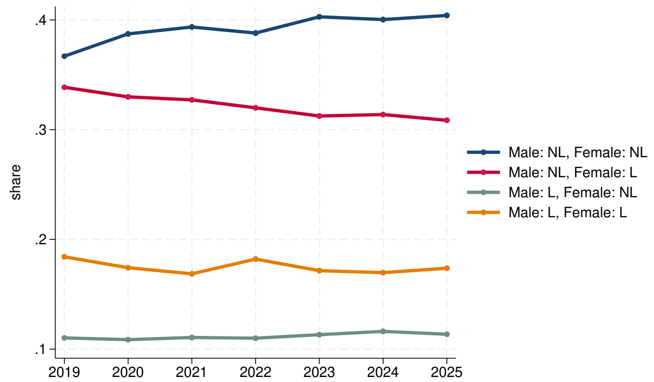

Note: As Figure 8, focusing on the window around 2020.

Figure A.4: Teleworkability and AI exposure: alternative measures

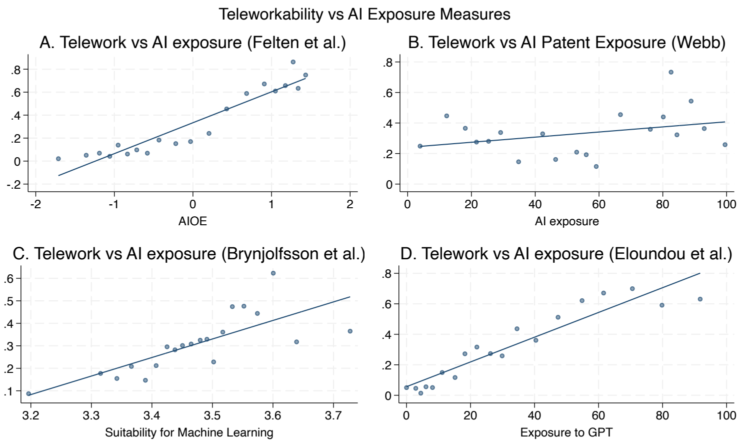

Note: Teleworkability (Dingel and Neiman, 2020) against four AI exposure measures: AIOE (Felten, Raj, and Seamans, 2021), AI patent exposure (Webb, 2020), suitability for machine learning (Brynjolfsson, Mitchell, and Rock, 2018), and exposure to GPTs (Eloundou et al., 2024).

Figure A.5: AI exposure by hours rank and gender: robustness

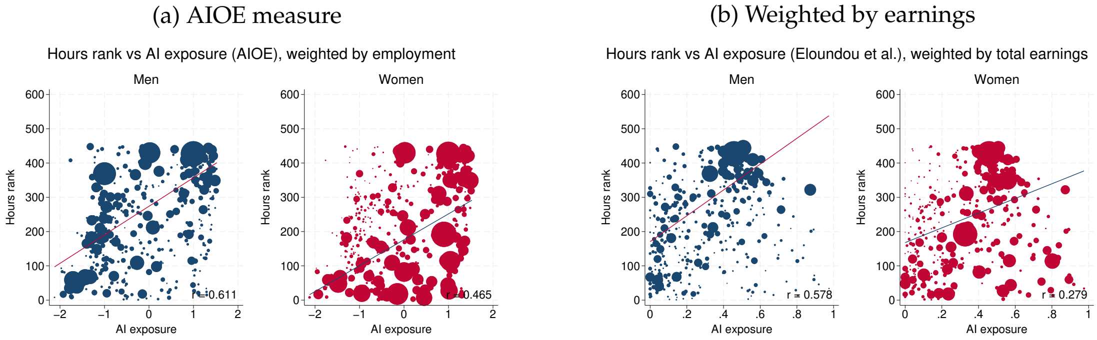

Note: As Panel (b) of Figure 17, using the Felten, Raj, and Seamans (2021) AIOE measure (panel a) and weighting occupations by earnings (panel b).

Figure A.6: Cumulative distribution of AI exposure by gender

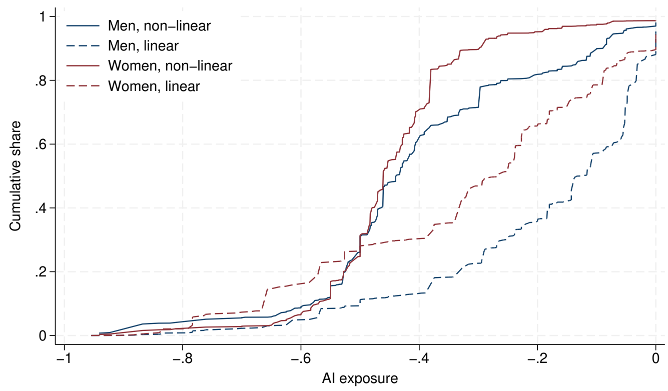

Note: Employment-weighted cumulative distribution of occupation-level AI exposure, by gender.

B From tasks to the aggregate production function

This appendix details the aggregation underlying the production function (15), which follows Acemoglu and Restrepo (2018), and the mapping from measured AI exposure into the model. The material is standard; we record it to fix notation, to make explicit how differences between men and women are treated, and to connect the model objects to the data moments used to define AI exposure.

Tasks and comparative advantage. Output is produced by combining a unit continuum of tasks $x \in [0, 1]$. Each task can be performed by capital, by unskilled labor (supplied by workers in the linear occupation), or by skilled labor (workers in the non-linear occupation). Comparative advantage is log-linear in the task index: the productivity of capital relative to unskilled labor declines in x at rate χ, and that of unskilled relative to skilled labor declines at rate η. Cost minimization then implies two thresholds, I and $\tilde{x}$, such that capital performs tasks $[0, I]$, unskilled labor $(I, \tilde{x}]$, and skilled labor $(\tilde{x}, 1]$. The thresholds equate the effective unit cost of adjacent factors: at I the relative productivity of capital equals the ratio of the unskilled wage to the user cost of capital, and at $\tilde{x}$ the relative productivity schedule equals the skill premium w_S/w_U.

Aggregation. Integrating task-level outputs across the three task ranges delivers the CES representation in the main text,

$$Y = A\left[m_K\,(A_K K)^{\zeta_c} + m_U\,(A_U U)^{\zeta_c} + m_S\,(A_S S)^{\zeta_c}\right]^{1/\zeta_c}, \tag{A.1}$$

in which the share parameters equal the measures of tasks assigned to each factor. Absent automation, $m_K = I$, $m_U = \tilde{x} - I$, and $m_S = 1 - \tilde{x}$. When AI automates a fraction a_U of the tasks currently performed by unskilled labor and a fraction a_S of those performed by skilled labor, the task measures become

$$m_K = I + a_U(\tilde{x} - I) + a_S(1 - \tilde{x}), \qquad m_U = (1 - a_U)(\tilde{x} - I), \qquad m_S = (1 - a_S)(1 - \tilde{x}). \tag{A.2}$$

The cost saving on automated tasks raises aggregate productivity: output and all factor prices are scaled by the Hulten term associated with the measure of automated tasks, so that the productivity gain from automation accrues to labor as well as to capital (Acemoglu, 2025). Absent automation the scaling equals one and the framework reduces to the three-factor CES used in the work-from-home experiments.

Calibration. At the initial steady state there is no automation and per-worker wages are normalized to one, so the thresholds follow from Cobb–Douglas share identities: $I_0 = r_K K/Y_0$ and $\tilde{x}_0 = I_0 + U/Y_0$, where r_K is the user cost of capital and $Y_0 = r_K K + U + S$. The slope of the unskilled–skilled comparative advantage schedule is then set so that the local elasticity of substitution between the two occupations at the initial thresholds equals the target value,

$$\eta = \frac{1}{(\epsilon_L - 1)\,\tilde{x}_0\,(1 - \tilde{x}_0)}, \qquad \epsilon_L = 1.6, \tag{A.3}$$

so that for small shocks the framework reproduces the CES benchmark of Autor, Katz, and Kearney (2008). The rental rate of capital is held at its baseline value, with the capital stock adjusting.

Treatment of gender. Men and women are perfect substitutes within a skill cell, weighted by group-specific productivities:

$$S = \phi_S^m\, e_S^m + \phi_S^f\, e_S^f, \qquad U = \phi_U^m\, e_U^m + \phi_U^f\, e_U^f, \tag{A.4}$$

where e_o^g denotes the aggregate efficiency units supplied by gender g in occupation o and $\phi_o^g = 1$ in the absence of AI. The firm pays a single price per effective unit in each cell (w_U, w_S); the wage received by group g per efficiency unit supplied is ϕ_o^g times the cell price. Differential AI exposure across genders therefore maps one-for-one into within-cell relative wages, without introducing any additional substitution parameter, and the firm's task-assignment problem depends on gender composition only through the cell-level aggregates defined below. When exposures are equal across genders, the model reduces exactly to the canonical one-gender Acemoglu–Restrepo framework.

Mapping measured AI exposure into the model. We take task-level exposure estimates from Eloundou et al. (2024) and aggregate them to the model's occupation-by-gender cells using the gender-specific composition of employment across detailed occupations within each of the two occupation groups. This yields, for each cell (o, g), the share of the cell's tasks exposed to AI; in the baseline experiment we assume that one half of exposed tasks is ultimately automated, defining the automated task share Δ_o^g and hence the post-AI productivity $\phi_o^g = 1 - \Delta_o^g$. On the firm side, the automation measures entering the task thresholds and the aggregate productivity term are the labor-weighted cell averages,

$$a_o = \frac{\sum_g \Delta_o^g L_o^g}{\sum_g L_o^g}, \qquad o \in \{U, S\}, \tag{A.5}$$

which drive I, $\tilde{x}$, and the Hulten term exactly as in the canonical model, while the group productivities ϕ_o^g allocate the incidence of automation across genders within each cell. The gendered exposure pattern documented in Panel (b) of Figure 17—similar exposure of male and female non-linear occupations, but roughly twice the exposure for women within linear occupations—thus enters the model as $\Delta_U^f \approx 2\,\Delta_U^m$ together with $\Delta_S^f \approx \Delta_S^m$.

C Calibration and computation details

C.1 The work-from-home parameter

The parameter Φ_{WFH} is chosen so that the model's short-run response replicates the quasi-experimental evidence discussed in Section 4.4. The procedure holds the economy at its baseline stationary equilibrium: prices, marital matches, occupational assignments, and the joint distribution of households over states are all fixed at their baseline values, so that the response reflects only labor supply decisions, as in the empirical designs. Given a candidate value of Φ, we re-solve the household decision rules under the new childcare cost and simulate the economy one year forward. We then compute the change in the model's motherhood employment gap—the employment rate of women with a young child relative to comparable women without children, conditional on age—relative to its baseline value of -11.9 percentage points, and bisect on Φ until the one-year change equals the 2.27 percentage point aggregate target. This delivers $\Phi_{\text{WFH}} = 0.4$.

As an alternative that mirrors the event-study designs of Harrington and Kahn (2025) and Basso et al. (2026) more closely, we also implement the target in a simulated stochastic panel: we track households in event time around the birth of the first child and compute a difference-in-differences of mothers' employment, comparing cohorts whose births occur before and after the introduction of WFH, at the post-birth horizon used in the empirical studies. Targeting the ≈ 2 percentage point causal estimate in this design delivers a very similar value of Φ.

For the asymmetric experiment reported in Table A.1, a single scale parameter is bisected on the same target while the proportional cost reduction in the non-linear occupation is scaled up by the relative incidence of remote work in non-linear versus linear occupations, $1 - \Phi_{NL} = \kappa\,(1 - \Phi_L)$ with $\kappa \approx 1.24$, yielding $(\Phi_{NL}, \Phi_L) = (0.35, 0.47)$.

C.2 The occupational taste dispersion

The dispersion of the initial occupational taste shock, σ_o, governs the elasticity of entrants' occupational choices to occupational values, and is disciplined by the estimate in Traiberman (2019) that a 1 percent increase in an occupation's wage raises the probability of choosing that occupation by 4.4 percentage points. We implement the target as follows. We perturb the wage per efficiency unit in the non-linear occupation, ω_{NL}, by 1 percent; re-solve the life-cycle problem and recompute the ex-ante occupational values $\hat{V}^{o_x, o_{x'}}$ of equation (11), holding the entry costs $\Omega(o_x)$ fixed at the values that rationalize the baseline occupational distribution; and evaluate the change in the share of entrants choosing the non-linear occupation implied by the Gumbel choice probabilities. Because participation, hours, and worker composition also respond, a 1 percent perturbation of the wage *rate* does not move observed average wages one-for-one; we therefore simulate the baseline and perturbed economies and rescale the choice response by the realized change in the employment-weighted average non-linear wage, so that the calibrated statistic is the percentage-point change in the entry share per 1 percent change in *observed* average occupational wages—the empirical object in Traiberman (2019). Setting this statistic to 4.4 yields $\sigma_o = 1.30$; perturbing the male wage alone or the wages of both genders gives similar values.

C.3 The transition path

The transition paths in Figure 16 are computed as follows. At date one the work-from-home parameter changes permanently and unexpectedly; thereafter households have perfect foresight over the price path. We approximate the equilibrium price path by a step function: a constant vector of transition prices—one wage per efficiency unit for each occupation, common across genders—governs the initial phase of the transition, after which prices take their new steady-state values.

Given candidate transition prices, we (i) solve the household problem in the transition price environment; (ii) roll the pre-shock cohorts forward from the baseline stationary distribution under the new decision rules, so that incumbents adjust only through labor supply and the frictional occupational switching margin; (iii) let each entering cohort re-optimize its initial occupational choice under transition prices, holding the entry costs at their calibrated values; and (iv) aggregate efficiency units of labor by occupation period by period, combining the simulated paths of pre-shock cohorts and transition entrants with the new steady state to which the economy converges. We then compute the wages implied by the production function along this aggregate path and update the candidate transition prices until the prices households face coincide, on average over the first decade of the transition, with the implied prices. When capital enters production, the interest rate along the transition is determined by the same fixed point.

D Additional quantitative results

Asymmetric work-from-home shock. The baseline experiment applies a uniform reduction in the childcare cost scale Φ. Because non-linear occupations are more teleworkable (Figure 5), we also consider an experiment in which each spouse's flexibility gain scales with the teleworkability of his or her occupation, implying a WFH effect approximately 25 percent larger in the non-linear sector, $(\Phi_{NL}, \Phi_L) = (0.35, 0.47)$; Appendix C.1 details the mapping. Table A.1 reports the results. The asymmetric shock induces a larger shift of female labor supply toward non-linear occupations (12.8 versus 10.9 percentage points in general equilibrium) and a correspondingly larger general equilibrium wage response, which further amplifies male reallocation toward linear occupations. Aggregate income and employment responses are close to the baseline experiment.

Table A.1: Aggregate responses: teleworkability-scaled (asymmetric) WFH shock

Change	Male			Female		
	5 years	P.E.	G.E.	5 years	P.E.	G.E.
Income (%)	−0.4	−4.6	−7.6	4.2	14.7	14.1
Employment (p.p.)	−0.5	−2.2	−1.9	1.7	0.8	1.7
Employment with young child (p.p.)	−3.2	−4.1	−4.0	9.1	6.7	7.7
Hours (%)	0.0	−2.0	−3.8	2.6	9.6	8.4
Wages (%)	0.4	0.6	−0.6	−0.2	4.4	3.8
Non-linear share (p.p.)	−0.3	−6.5	−10.3	1.7	14.6	12.8

Note: As Table 6, for the experiment in which the WFH reduction in childcare costs scales with the teleworkability of each spouse's occupation.